



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

| | | |
|---|-----------|---|
| (51) International Patent Classification ⁶ : G06K 9/36 | A1 | (11) International Publication Number: WO 97/43736 (43) International Publication Date: 20 November 1997 (20.11.97) |
|---|-----------|---|

(21) International Application Number: PCT/US97/08351

(22) International Filing Date: 16 May 1997 (16.05.97)

(30) Priority Data:

| | | |
|------------|-----------------------------|----|
| 08/649,419 | 16 May 1996 (16.05.96) | US |
| 08/746,613 | 12 November 1996 (12.11.96) | US |

(60) Parent Applications or Grants

(63) Related by Continuation

| | |
|----------|-----------------------------|
| US | 08/649,419 (CIP) |
| Filed on | 16 May 1996 (16.05.96) |
| US | 08/746,613 (CIP) |
| Filed on | 12 November 1996 (12.11.96) |

(71) Applicant (for all designated States except US): DIGIMARC CORPORATION [US/US]; 521 S.W. 11th Avenue #200, Portland, OR 97205 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): RHOADS, Geoffrey, B. [US/US]; 304 S.W. Tualatin Loop, West Linn, OR 97068 (US).


(74) Agents: JOHNSON, Alexander, C., Jr. et al.; Marger, Johnson, McCollom & Stolowitz, 1030 S.W. Morrison Street, Portland, OR 97205 (US).

(81) Designated States: AU, CA, JP, US, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published

*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(54) Title: COMPUTER SYSTEM LINKED BY USING INFORMATION IN DATA OBJECTS



(57) Abstract

Various improvements to steganographic systems, and applications therefore, are disclosed. The improvements include facilitating scale and rotation registration for steganographic decoding by use of rotationally symmetric steganographically embedded patterns and subliminal digital gratitudes; improved techniques for decoding without access to unencoded originals; improving robustness of steganographic coding in motion pictures and/or in the presence of lossy compression/decompression; and representing data (1690) by patterned bit cells (802) whose energy in the spatial domain facilitates decoding registration (1692). Applications include enhanced-security financial transactions, counterfeit resistant identification cards, fraud deterrent systems for cellular telephony, covert modem channels in video transmissions, photo duplication kiosks with automatic copyright detection, and hotlinked image objects (e.g., with embedded URLs) for use on the Internet.

BEST AVAILABLE COPY

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

| | | | | | | | |
|----|--------------------------|----|--|----|--|----|--------------------------|
| AL | Albania | ES | Spain | LS | Lesotho | SI | Slovenia |
| AM | Armenia | FI | Finland | LT | Lithuania | SK | Slovakia |
| AT | Austria | FR | France | LU | Luxembourg | SN | Senegal |
| AU | Australia | GA | Gabon | LV | Latvia | SZ | Swaziland |
| AZ | Azerbaijan | GB | United Kingdom | MC | Monaco | TD | Chad |
| BA | Bosnia and Herzegovina | GE | Georgia | MD | Republic of Moldova | TG | Togo |
| BB | Barbados | GH | Ghana | MG | Madagascar | TJ | Tajikistan |
| BE | Belgium | GN | Guinea | MK | The former Yugoslav Republic of Macedonia | TM | Turkmenistan |
| BF | Burkina Faso | GR | Greece | ML | Mali | TR | Turkey |
| BG | Bulgaria | HU | Hungary | MN | Mongolia | TT | Trinidad and Tobago |
| BJ | Benin | IE | Ireland | MR | Mauritania | UA | Ukraine |
| BR | Brazil | IL | Israel | MW | Malawi | UG | Uganda |
| BY | Belarus | IS | Iceland | MX | Mexico | US | United States of America |
| CA | Canada | IT | Italy | NE | Niger | UZ | Uzbekistan |
| CF | Central African Republic | JP | Japan | NL | Netherlands | VN | Viet Nam |
| CG | Congo | KE | Kenya | NO | Norway | YU | Yugoslavia |
| CH | Switzerland | KG | Kyrgyzstan | NZ | New Zealand | ZW | Zimbabwe |
| CI | Côte d'Ivoire | KP | Democratic People's Republic of Korea | PL | Poland | | |
| CM | Cameroon | KR | Republic of Korea | PT | Portugal | | |
| CN | China | KZ | Kazakstan | RO | Romania | | |
| CU | Cuba | LC | Saint Lucia | RU | Russian Federation | | |
| CZ | Czech Republic | LI | Liechtenstein | SD | Sudan | | |
| DE | Germany | LK | Sri Lanka | SE | Sweden | | |
| DK | Denmark | LR | Liberia | SG | Singapore | | |
| EE | Estonia | | | | | | |

COMPUTER SYSTEM LINKED BY USING
INFORMATION IN DATA OBJECTS

Background of the Invention

Hiding data in imagery or audio is a technique well known to artisans in the field,
5 and is termed "steganography." There are a number of diverse approaches to, and
applications of, steganography. A brief survey follows:

British patent publication 2,196,167 to Thorn EMI discloses a system in which an
audio recording is electronically mixed with a marking signal indicative of the owner of
the recording, where the combination is perceptually identical to the original. U.S. patents
10 4,963,998 and 5,079,648 disclose variants of this system.

U.S. Patent 5,319,735 to Bolt, Berenak & Newman rests on the same principles as
the earlier Thorn EMI publication, but additionally addresses psycho-acoustic masking
issues.

U.S. Patents 4,425,642, 4,425,661, 5,404,377 and 5,473,631 to Moses disclose
15 various systems for imperceptibly embedding data into audio signals -- the latter two
patents particularly focusing on neural network implementations and perceptual coding
details.

U.S. Patent 4,943,973 to AT&T discloses a system employing spread spectrum
techniques for adding a low level noise signal to other data to convey auxiliary data
20 therewith. The patent is particularly illustrated in the context of transmitting network
control signals along with digitized voice signals.

U.S. Patent 5,161,210 to U.S. Philips discloses a system in which additional low-
level quantization levels are defined on an audio signal to convey, e.g., a copy inhibit code,
therewith.

25 U.S. Patent 4,972,471 to Gross discloses a system intended to assist in the
automated monitoring of audio (e.g. radio) signals for copyrighted materials by reference
to identification signals subliminally embedded therein.

U.S. Patent 5,243,423 to DeJean discloses a video steganography system which
encodes digital data (e.g. program syndication verification, copyright marking, media
30 research, closed captioning, or like data) onto randomly selected video lines. DeJean relies

on television sync pulses to trigger a stored pseudo random sequence which is XORed with the digital data and combined with the video.

European application EP 581,317 discloses a system for redundantly marking images with multi-bit identification codes. Each "1" ("0") bit of the code is manifested as a slight increase (decrease) in pixel values around a plurality of spaced apart "signature points." Decoding proceeds by computing a difference between a suspect image and the original, unencoded image, and checking for pixel perturbations around the signature points.

PCT application WO 95/14289 describes the present applicant's prior work in this field.

Komatsu et al., describe an image marking technique in their paper "A Proposal on Digital Watermark in Document Image Communication and Its Application to Realizing a Signature," Electronics and Communications in Japan, Part 1, Vol. 73, No. 5, 1990, pp. 22-33. The work is somewhat difficult to follow but apparently results in a simple yes/no determination of whether the watermark is present in a suspect image (e.g. a 1 bit encoded message).

Others embed the information "in-band" (i.e. in the visible video signal itself). Examples include U.S. Patents 4,528,588, 4,595,950, and 5,319,453; European application 441,702; and Matsui et. al, "Video-Steganography: How to Secretly Embed a Signature in a Picture," IMA Intellectual Property Project Proceedings, January 1994, Vol. 1, Issue 1, pp. 187-205.

Copyright marking of video and multimedia is found in "Access Control and Copyright Protection for Images (ACCOPI), WorkPackage 8: Watermarking," June 30, 1995, 46 pages. TALISMAN, appears to extend certain of the ACCOPI work.

A British company, Highwater FBI, Ltd., has introduced a software product which is said to imperceptibly embed identifying information into photographs and other graphical images. This technology is the subject of European patent applications 9400971.9 (filed January 19, 1994), 9504221.2 (filed March 2, 1995), and 9513790.7 (filed July 3, 1995), the first of which has been laid open as PCT publication WO 95/20291.

Brief Description Of The Drawings

Fig. 1 is a simple and classic depiction of a one dimensional digital signal which is discretized in both axes.

5 Fig. 2 is a general overview, with detailed description of steps, of the process of embedding an "imperceptible" identification signal onto another signal.

Fig. 3 is a step-wise description of how a suspected copy of an original is identified.

Fig. 4 is a schematic view of an apparatus for pre-exposing film with identification information.

Fig. 5 is a diagram of a "black box" embodiment.

10 Fig. 6 is a schematic block diagram of the embodiment of Fig. 5.

Fig. 7 shows a variant of the Fig. 6 embodiment adapted to encode successive sets of input data with different code words but with the same noise data.

Fig. 8 shows a variant of the Fig. 6 embodiment adapted to encode each frame of a videotaped production with a unique code number.

15 Figs. 9A-9C are representations of an industry standard noise second.

Fig. 10 shows an integrated circuit used in detecting standard noise codes.

Fig. 11 shows a process flow for detecting a standard noise code that can be used in the Fig. 10 embodiment.

Fig. 12 is an embodiment employing a plurality of detectors.

20 Fig. 13 shows an embodiment in which a pseudo-random noise frame is generated from an image.

Fig. 14 illustrates how statistics of a signal can be used in aid of decoding.

Fig. 15 shows how a signature signal can be preprocessed to increase its robustness in view of anticipated distortion, e.g. MPEG.

25 Figs. 16 and 17 show embodiments in which information about a file is detailed both in a header, and in the file itself.

Figs. 18-20 show details of embodiments using rotationally symmetric patterns.

Fig. 21 shows encoding "bumps" rather than pixels.

Figs. 22-26 detail aspects of a security card.

30 Fig. 27 is a diagram illustrating a network linking method using information embedded in data objects that have inherent noise.

Figs. 27A and 27B show a typical web page, and a step in its encapsulation into a self extracting web page object.

Fig. 28 is a diagram of a photographic identification document or security card.

5 Figs. 29 and 30 illustrate two embodiments by which subliminal digital graticules can be realized.

Fig. 29A shows a variation on the Fig. 29 embodiment.

Figs. 31A and 31B show the phase of spatial frequencies along two inclined axes.

10 Figs. 32A - 32C show the phase of spatial frequencies along first, second and third concentric rings.

Figs 33A - 33E show steps in the registration process for a subliminal graticule using inclined axes.

Figs. 34A - 34E show steps in the registration process for a subliminal graticule using concentric rings.

15 Figs. 35A - 35C shows further steps in the registration process for a subliminal graticule using inclined axes.

Figs. 36A - 36D show another registration process that does not require a 2D FFT.

Fig. 37 is a flow chart summarizing a registration process for subliminal graticules.

Fig. 38 is a block diagram showing principal components of an exemplary wireless telephony system.

20 Fig. 39 is a block diagram of an exemplary steganographic encoder that can be used in the telephone of the Fig. 38 system.

Fig. 40 is a block diagram of an exemplary steganographic decoder that can be used in the cell site of the Fig. 1 system.

Figs. 41A and 41B show exemplary bit cells used in one form of encoding.

25 Fig. 42 shows a hierarchical arrangement of signature blocks, sub-blocks, and bit cells used in one embodiment.

Fig. 43 shows a general overview of a computer system linked by using information embedded in the objects.

30 Figs. 44-59 are representative computer screens generated in accordance with the present invention.

Fig. 60 is a basic overview of how a traditional file header can be replaced by a pointer to a database location where header information resides.

Fig. 61 expands on Fig. 60 by illustrating that a pointer to a remote location can coexist with existing data file formats which might include headers.

5 Fig. 62 is a basic map showing how an instrument or application can associate a pointer attached to an object to some remote digital storage location.

Fig. 63 is a condensed depiction of a process that an instrument or application might go through in resolving and acting upon remote information pointed to by an attached pointer.

10 Fig. 64 graphically illustrates that attached information can actually be sequences of actions or multiple sets of information.

Fig. 65 shows how the basic principles of master identity objects can be used to fuel a reporting network wherein a central location can "keep track" of where its copies migrate.

15 Detailed Description

In the following discussion of an illustrative embodiment, the words "signal" and "image" are used interchangeably to refer to both one, two, and even beyond two dimensions of digital signal. Examples will routinely switch back and forth between a one dimensional audio-type digital signal and a two dimensional image-type
20 digital signal.

In order to fully describe the details of an illustrative embodiment, it is necessary first to describe the basic properties of a digital signal. Fig. 1 shows a classic representation of a one dimensional digital signal. The x-axis defines the index numbers of sequence of digital "samples," and the y-axis is the instantaneous
25 value of the signal at that sample, being constrained to exist only at a finite number of levels defined as the "binary depth" of a digital sample. The example depicted in Fig. 1 has the value of 2 to the fourth power, or "4 bits," giving 16 allowed states of the sample value.

- 6 -

For audio information such as sound waves, it is commonly accepted that the digitization process discretizes a continuous phenomena both in the time domain and in the signal level domain. As such, the process of digitization itself introduces a fundamental error source, in that it cannot record detail smaller than the discretization interval in either domain. The industry has referred to this, among other ways, as "aliasing" in the time domain, and "quantization noise" in the signal level domain. Thus, there will always be a basic error floor of a digital signal. Pure quantization noise, measured in a root mean square sense, is theoretically known to have the value of one over the square root of twelve, or about 0.29 DN, where DN stands for 'Digital Number' or the finest unit increment of the signal level. For example, a perfect 12-bit digitizer will have 4096 allowed DN with an innate root mean square noise floor of ~ 0.29 DN.

10 All known physical measurement processes add additional noise to the transformation of a continuous signal into the digital form. The quantization noise typically adds in quadrature (square root of the mean squares) to the "analog noise" of the measurement process, as it is sometimes referred to.

With almost all commercial and technical processes, the use of the decibel scale is used as a measure of signal and noise in a given recording medium. The expression "signal-to-noise ratio" is generally used, as it will be in this disclosure. As an example, this disclosure refers to signal to noise ratios in terms of signal power and noise power, thus 20 dB represents a 10 times increase in signal amplitude.

In summary, this embodiment embeds an N-bit value onto an entire signal through the addition of a very low amplitude encodation signal which has the look of pure noise. N is usually at least 8 and is capped on the higher end by ultimate signal-to-noise considerations and "bit error" in retrieving and decoding the N-bit value. As a practical matter, N is chosen based on application specific considerations, such as the number of unique different "signatures" that are desired. To illustrate, if $N=128$, then the number of unique digital signatures is in excess of 10^{38} (2^{128}). This number is believed to be more than adequate to both identify the material with sufficient statistical certainty and to index exact sale and distribution information.

25 The amplitude or power of this added signal is determined by the aesthetic and informational considerations of each and every application using the present methodology. For instance, non-professional video can stand to have a higher embedded signal level without becoming noticeable to the average human eye, while high precision audio may only be able to accept a relatively small signal level lest the human ear perceive an objectionable increase in "hiss." These statements are generalities and each application has its own set of criteria in choosing the signal level of the embedded identification signal. The higher the level of embedded signal, the more corrupted a copy can be and still be identified. On the other hand, the higher the level of embedded signal, the more objectionable the perceived noise might be, potentially impacting the value of the distributed material.

To illustrate the range of different applications to which applicant's technology can be applied, the present specification details two different systems. The first (termed, for lack of a better name, a "batch encoding" system), applies identification coding to an existing data signal. The second (termed, for lack of a better name, a "real time encoding" system), applies identification coding to a signal as it is produced.

5 Those skilled in the art will recognize that the principles of applicant's technology can be applied in a number of other contexts in addition to these particularly described.

The discussions of these two systems can be read in either order. Some readers may find the latter more intuitive than the former; for others the contrary may be true.

10 BATCH ENCODING

The following discussion of a first class of embodiments is best prefaced by a section defining relevant terms:

The original signal refers to either the original digital signal or the high quality digitized copy of a non-digital original.

15 The N-bit identification word refers to a unique identification binary value, typically having N range anywhere from 8 to 128, which is the identification code ultimately placed onto the original signal via the disclosed transformation process. In the illustrated embodiment, each N-bit identification word begins with the sequence of values '0101,' which is used to determine an optimization of the signal-to-noise ratio in the identification procedure of a suspect signal (see definition below).

20 The m'th bit value of the N-bit identification word is either a zero or one corresponding to the value of the m'th place, reading left to right, of the N-bit word. E.g., the first (m=1) bit value of the N=8 identification word 01110100 is the value '0;' the second bit value of this identification word is '1,' etc.

25 The m'th individual embedded code signal refers to a signal which has dimensions and extent precisely equal to the original signal (e.g. both are a 512 by 512 digital image), and which is (in the illustrated embodiment) an independent pseudo-random sequence of digital values. "Pseudo" pays homage to the difficulty in philosophically defining pure randomness, and also indicates that there are various acceptable ways of generating the "random" signal. There will be exactly N individual embedded code signals associated with any given original signal.

30 The acceptable perceived noise level refers to an application-specific determination of how much "extra noise," i.e. amplitude of the composite embedded code signal described next, can be added to the original signal and still have an acceptable signal to sell or otherwise distribute. This disclosure uses a 1 dB increase in noise as a typical value which might be acceptable, but this is quite arbitrary.

35 The composite embedded code signal refers to the signal which has dimensions and extent precisely equal to the original signal, (e.g. both are a 512 by 512 digital image), and which contains the

addition and appropriate attenuation of the N individual embedded code signals. The individual embedded signals are generated on an arbitrary scale, whereas the amplitude of the composite signal must not exceed the pre-set acceptable perceived noise level, hence the need for "attenuation" of the N added individual code signals.

5 The distributable signal refers to the nearly similar copy of the original signal, consisting of the original signal plus the composite embedded code signal. This is the signal which is distributed to the outside community, having only slightly higher but acceptable "noise properties" than the original.

10 A suspect signal refers to a signal which has the general appearance of the original and distributed signal and whose potential identification match to the original is being questioned. The suspect signal is then analyzed to see if it matches the N-bit identification word.

15 The detailed methodology of this first embodiment begins by stating that the N-bit identification word is encoded onto the original signal by having each of the m bit values multiply their corresponding individual embedded code signals, the resultant being accumulated in the composite signal, the fully summed composite signal then being attenuated down to the acceptable perceived noise amplitude, and the resultant composite signal added to the original to become the distributable signal.

20 The original signal, the N-bit identification word, and all N individual embedded code signals are then stored away in a secured place. A suspect signal is then found. This signal may have undergone multiple copies, compressions and decompressions, resamplings onto different spaced digital signals, transfers from digital to analog back to digital media, or any combination of these items. IF the signal still appears similar to the original, i.e. its innate quality is not thoroughly destroyed by all of these transformations and noise additions, then depending on the signal to noise properties of the embedded signal, the identification process should function to some objective degree of statistical confidence. The extent of corruption of the suspect signal and the original acceptable perceived noise level are two key parameters in determining an expected confidence level of identification.

25 The identification process on the suspected signal begins by resampling and aligning the suspected signal onto the digital format and extent of the original signal. Thus, if an image has been reduced by a factor of two, it needs to be digitally enlarged by that same factor. Likewise, if a piece of music has been "cut out," but may still have the same sampling rate as the original, it is necessary to register this cut-out piece to the original, typically done by performing a local digital cross-correlation of the two signals (a common digital operation), finding at what delay value the correlation peaks, then using this found delay value to register the cut piece to a segment of the original.

30 Once the suspect signal has been sample-spacing matched and registered to the original, the signal levels of the suspect signal should be matched in an rms sense to the signal level of the original. This can be done via a search on the parameters of offset, amplification, and gamma being optimized by using the

- 9 -

minimum of the mean squared error between the two signals as a function of the three parameters. We can call the suspect signal normalized and registered at this point, or just normalized for convenience.

The newly matched pair then has the original signal subtracted from the normalized suspect signal to produce a difference signal. The difference signal is then cross-correlated with each of the N individual
5 embedded code signals and the peak cross-correlation value recorded. The first four bit code ('0101') is used as a calibrator both on the mean values of the zero value and the one value, and on further registration of the two signals if a finer signal to noise ratio is desired (i.e., the optimal separation of the 0101 signal will indicate an optimal registration of the two signals and will also indicate the probable existence of the N -bit identification signal being present.)

10 The resulting peak cross-correlation values will form a noisy series of floating point numbers which can be transformed into 0's and 1's by their proximity to the mean values of 0 and 1 found by the 0101 calibration sequence. If the suspect signal has indeed been derived from the original, the identification number resulting from the above process will match the N -bit identification word of the original, bearing in mind either predicted or unknown "bit error" statistics. Signal-to-noise considerations
15 will determine if there will be some kind of "bit error" in the identification process, leading to a form of $X\%$ probability of identification where X might be desired to be 99.9% or whatever. If the suspect copy is indeed not a copy of the original, an essentially random sequence of 0's and 1's will be produced, as well as an apparent lack of separation of the resultant values. This is to say, if the resultant values are plotted on a histogram, the existence of the N -bit identification signal will exhibit strong bi-level
20 characteristics, whereas the non-existence of the code, or the existence of a different code of a different original, will exhibit a type of random gaussian-like distribution. This histogram separation alone should be sufficient for an identification, but it is even stronger proof of identification when an exact binary sequence can be objectively reproduced.

25 Specific Example

Imagine that we have taken a valuable picture of two heads of state at a cocktail party, pictures which are sure to earn some reasonable fee in the commercial market. We desire to sell this picture and ensure that it is not used in an unauthorized or uncompensated manner. This and the following steps are summarized in Fig. 2.

30 Assume the picture is transformed into a positive color print. We first scan this into a digitized form via a normal high quality black and white scanner with a typical photometric spectral response curve. (It is possible to get better ultimate signal to noise ratios by scanning in each of the three primary colors of the color image, but this nuance is not central to describing the basic process.)

- 10 -

Let us assume that the scanned image now becomes a 4000 by 4000 pixel monochrome digital image with a grey scale accuracy defined by 12-bit grey values or 4096 allowed levels. We will call this the "original digital image" realizing that this is the same as our "original signal" in the above definitions.

During the scanning process we have arbitrarily set absolute black to correspond to digital value '30'. We estimate that there is a basic 2 Digital Number root mean square noise existing on the original digital image, plus a theoretical noise (known in the industry as "shot noise") of the square root of the brightness value of any given pixel. In formula, we have:

$$\langle \text{RMS Noise}_{n,m} \rangle = \text{sqrt}(4 + (V_{n,m} - 30)) \quad (1)$$

10

Here, n and m are simple indexing values on rows and columns of the image ranging from 0 to 3999. Sqrt is the square root. V is the DN of a given indexed pixel on the original digital image. The < > brackets around the RMS noise merely indicates that this is an expected average value, where it is clear that each and every pixel will have a random error individually. Thus, for a pixel value having 1200 as a digital number or "brightness value", we find that its expected rms noise value is $\text{sqrt}(1204) = 34.70$, which is quite close to 34.64, the square root of 1200.

15

We furthermore realize that the square root of the innate brightness value of a pixel is not precisely what the eye perceives as a minimum objectionable noise, thus we come up with the formula:

20

$$\langle \text{RMS Addable Noise}_{n,m} \rangle = X * \text{sqrt}(4 + (V_{n,m} - 30)^Y) \quad (2)$$

25

Where X and Y have been added as empirical parameters which we will adjust, and "addable" noise refers to our acceptable perceived noise level from the definitions above. We now intend to experiment with what exact value of X and Y we can choose, but we will do so at the same time that we are performing the next steps in the process.

30

The next step in our process is to choose N of our N-bit identification word. We decide that a 16 bit main identification value with its 65536 possible values will be sufficiently large to identify the image as ours, and that we will be directly selling no more than 128 copies of the image which we wish to track, giving 7 bits plus an eighth bit for an odd/even adding of the first 7 bits (i.e. an error checking bit on the first seven). The total bits required now are at 4 bits for the 0101 calibration sequence, 16 for the main identification, 8 for the version, and we now throw in another 4 as a further error checking value on the first 28 bits, giving 32 bits as N. The final 4 bits can use one of many industry standard error checking methods to choose its four values.

35

We now randomly determine the 16 bit main identification number, finding for example, 1101 0001 1001 1110; our first versions of the original sold will have all 0's as the version identifier, and the

- 11 -

error checking bits will fall out where they may. We now have our unique 32 bit identification word which we will embed on the original digital image.

To do this, we generate 32 independent random 4000 by 4000 encoding images for each bit of our 32 bit identification word. The manner of generating these random images is revealing. There are numerous ways to generate these. By far the simplest is to turn up the gain on the same scanner that was used to scan in the original photograph, only this time placing a pure black image as the input, then scanning this 32 times. The only drawback to this technique is that it does require a large amount of memory and that "fixed pattern" noise will be part of each independent "noise image." But, the fixed pattern noise can be removed via normal "dark frame" subtraction techniques. Assume that we set the absolute black average value at digital number '100,' and that rather than finding a 2 DN rms noise as we did in the normal gain setting, we now find an rms noise of 10 DN about each and every pixel's mean value.

We next apply a mid-spatial-frequency bandpass filter (spatial convolution) to each and every independent random image, essentially removing the very high and the very low spatial frequencies from them. We remove the very low frequencies because simple real-world error sources like geometrical warping, splotches on scanners, mis-registrations, and the like will exhibit themselves most at lower frequencies also, and so we want to concentrate our identification signal at higher spatial frequencies in order to avoid these types of corruptions. Likewise, we remove the higher frequencies because multiple generation copies of a given image, as well as compression-decompression transformations, tend to wipe out higher frequencies anyway, so there is no point in placing too much identification signal into these frequencies if they will be the ones most prone to being attenuated. Therefore, our new filtered independent noise images will be dominated by mid-spatial frequencies. On a practical note, since we are using 12-bit values on our scanner and we have removed the DC value effectively and our new rms noise will be slightly less than 10 digital numbers, it is useful to boil this down to a 6-bit value ranging from -32 through 0 to 31 as the resultant random image.

Next we add all of the random images together which have a '1' in their corresponding bit value of the 32-bit identification word, accumulating the result in a 16-bit signed integer image. This is the unattenuated and un-scaled version of the composite embedded signal.

Next we experiment visually with adding the composite embedded signal to the original digital image, through varying the X and Y parameters of equation 2. In formula, we visually iterate to both maximize X and to find the appropriate Y in the following:

$$V_{\text{dist},n,m} = V_{\text{orig},n,m} + V_{\text{comp},n,m} * X * \sqrt{4 + V_{\text{orig},n,m}^Y} \quad (3)$$

- 12 -

where dist refers to the candidate distributable image, i.e. we are visually iterating to find what X and Y will give us an acceptable image; orig refers to the pixel value of the original image; and comp refers to the pixel value of the composite image. The n's and m's still index rows and columns of the image and indicate that this operation is done on all 4000 by 4000 pixels. The symbol V is the DN of a given pixel and a given image.

As an arbitrary assumption, now, we assume that our visual experimentation has found that the value of $X = 0.025$ and $Y = 0.6$ are acceptable values when comparing the original image with the candidate distributable image. This is to say, the distributable image with the "extra noise" is acceptably close to the original in an aesthetic sense. Note that since our individual random images had a random rms noise value around 10 DN, and that adding approximately 16 of these images together will increase the composite noise to around 40 DN, the X multiplication value of 0.025 will bring the added rms noise back to around 1 DN, or half the amplitude of our innate noise on the original. This is roughly a 1 dB gain in noise at the dark pixel values and correspondingly more at the brighter values modified by the Y value of 0.6.

So with these two values of X and Y, we now have constructed our first versions of a distributable copy of the original. Other versions will merely create a new composite signal and possibly change the X slightly if deemed necessary. We now lock up the original digital image along with the 32-bit identification word for each version, and the 32 independent random 4-bit images, waiting for our first case of a suspected piracy of our original. Storage wise, this is about 14 Megabytes for the original image and $32 * 0.5 \text{ bytes} * 16 \text{ million} = \sim 256 \text{ Megabytes}$ for the random individual encoded images. This is quite acceptable for a single valuable image. Some storage economy can be gained by simple lossless compression.

Finding a Suspected Piracy of our Image

We sell our image and several months later find our two heads of state in the exact poses we sold them in, seemingly cut and lifted out of our image and placed into another stylized background scene. This new "suspect" image is being printed in 100,000 copies of a given magazine issue, let us say. We now go about determining if a portion of our original image has indeed been used in an unauthorized manner. Fig. 3 summarizes the details.

The first step is to take an issue of the magazine, cut out the page with the image on it, then carefully but not too carefully cut out the two figures from the background image using ordinary scissors. If possible, we will cut out only one connected piece rather than the two figures separately. We paste this onto a black background and scan this into a digital form. Next we electronically flag or mask out the black background, which is easy to do by visual inspection.

- 13 -

We now procure the original digital image from our secured place along with the 32-bit identification word and the 32 individual embedded images. We place the original digital image onto our computer screen using standard image manipulation software, and we roughly cut along the same borders as our masked area of the suspect image, masking this image at the same time in roughly the same manner. The word 'roughly' is used since an exact cutting is not needed, it merely aids the identification statistics to get it reasonably close.

Next we rescale the masked suspect image to roughly match the size of our masked original digital image, that is, we digitally scale up or down the suspect image and roughly overlay it on the original image. Once we have performed this rough registration, we then throw the two images into an automated scaling and registration program. The program performs a search on the three parameters of x position, y position, and spatial scale, with the figure of merit being the mean squared error between the two images given any given scale variable and x and y offset. This is a fairly standard image processing methodology. Typically this would be done using generally smooth interpolation techniques and done to sub-pixel accuracy. The search method can be one of many, where the simplex method is a typical one.

Once the optimal scaling and x-y position variables are found, next comes another search on optimizing the black level, brightness gain, and gamma of the two images. Again, the figure of merit to be used is mean squared error, and again the simplex or other search methodologies can be used to optimize the three variables. After these three variables are optimized, we apply their corrections to the suspect image and align it to exactly the pixel spacing and masking of the original digital image and its mask. We can now call this the standard mask.

The next step is to subtract the original digital image from the newly normalized suspect image only within the standard mask region. This new image is called the difference image.

Then we step through all 32 individual random embedded images, doing a local cross-correlation between the masked difference image and the masked individual embedded image. 'Local' refers to the idea that one need only start correlating over an offset region of ± 1 pixels of offset between the nominal registration points of the two images found during the search procedures above. The peak correlation should be very close to the nominal registration point of 0,0 offset, and we can add the 3 by 3 correlation values together to give one grand correlation value for each of the 32 individual bits of our 32-bit identification word.

After doing this for all 32 bit places and their corresponding random images, we have a quasi-floating point sequence of 32 values. The first four values represent our calibration signal of 0101. We now take the mean of the first and third floating point value and call this floating point value '0,' and we take the mean of the second and the fourth value and call this floating point value '1.' We then step through all remaining 28 bit values and assign either a '0' or a '1' based simply on which mean value they are closer to. Stated simply, if the suspect image is indeed a copy of our original, the embedded 32-bit

- 14 -

resulting code should match that of our records, and if it is not a copy, we should get general randomness. The third and the fourth possibilities of 3) Is a copy but doesn't match identification number and 4) isn't a copy but does match are, in the case of 3), possible if the signal to noise ratio of the process has plummeted, i.e. the 'suspect image' is truly a very poor copy of the original, and in the case of 4) is basically one chance in four billion since we were using a 32-bit identification number. If we are truly worried about 4), we can just have a second independent lab perform their own tests on a different issue of the same magazine. Finally, checking the error-check bits against what the values give is one final and possibly overkill check on the whole process. In situations where signal to noise is a possible problem, these error checking bits might be eliminated without too much harm.

10

Benefits

Now that a full description of the first embodiment has been described via a detailed example, it is appropriate to point out the rationale of some of the process steps and their benefits.

The ultimate benefits of the foregoing process are that obtaining an identification number is fully independent of the manners and methods of preparing the difference image. That is to say, the manners of preparing the difference image, such as cutting, registering, scaling, etcetera, cannot increase the odds of finding an identification number when none exists; it only helps the signal-to-noise ratio of the identification process when a true identification number is present. Methods of preparing images for identification can be different from each other even, providing the possibility for multiple independent methodologies for making a match.

The ability to obtain a match even on sub-sets of the original signal or image is a key point in today's information-rich world. Cutting and pasting both images and sound clips is becoming more common, allowing such an embodiment to be used in detecting a copy even when original material has been thus corrupted. Finally, the signal to noise ratio of matching should begin to become difficult only when the copy material itself has been significantly altered either by noise or by significant distortion; both of these also will affect that copy's commercial value, so that trying to thwart the system can only be done at the expense of a huge decrease in commercial value.

An early conception of this technology was the case where only a single "snowy image" or random signal was added to an original image, i.e. the case where $N=1$. "Decoding" this signal would involve a subsequent mathematical analysis using (generally statistical) algorithms to make a judgment on the presence or absence of this signal. The reason this approach was abandoned as the preferred embodiment was that there was an inherent gray area in the certainty of detecting the presence or absence of the signal. By moving onward to a multitude of bit planes, i.e. $N > 1$, combined with simple pre-defined algorithms prescribing the manner of choosing between a "0" and a "1", the certainty question moved from the realm of expert statistical analysis into the realm of guessing a random binary event such

- 15 -

as a coin flip. This is seen as a powerful feature relative to the intuitive acceptance of this technology in both the courtroom and the marketplace. The analogy which summarizes the inventor's thoughts on this whole question is as follows: The search for a single identification signal amounts to calling a coin flip only once, and relying on arcane experts to make the call; whereas the $N > 1$ embodiment relies on the
5 broadly intuitive principle of correctly calling a coin flip N times in a row. This situation is greatly exacerbated, i.e. the problems of "interpretation" of the presence of a single signal, when images and sound clips get smaller and smaller in extent.

Another important reason that the $N > 1$ case is preferred over the $N = 1$ embodiment is that in the $N = 1$ case, the manner in which a suspect image is prepared and manipulated has a direct bearing on the
10 likelihood of making a positive identification. Thus, the manner with which an expert makes an identification determination becomes an integral part of that determination. The existence of a multitude of mathematical and statistical approaches to making this determination leave open the possibility that some tests might make positive identifications while others might make negative determinations, inviting further arcane debate about the relative merits of the various identification approaches. The $N > 1$ embodiment
15 avoids this further gray area by presenting a method where no amount of pre-processing of a signal - other than pre-processing which surreptitiously uses knowledge of the private code signals - can increase the likelihood of "calling the coin flip N times in a row."

The fullest expression of the present system will come when it becomes an industry standard and numerous independent groups set up with their own means or 'in-house' brand of applying embedded
20 identification numbers and in their decipherment. Numerous independent group identification will further enhance the ultimate objectivity of the method, thereby enhancing its appeal as an industry standard.

Use of True Polarity in Creating the Composite Embedded Code Signal

The foregoing discussion made use of the 0 and 1 formalism of binary technology to accomplish
25 its ends. Specifically, the 0's and 1's of the N -bit identification word directly multiplied their corresponding individual embedded code signal to form the composite embedded code signal (step 8, Fig. 2). This approach certainly has its conceptual simplicity, but the multiplication of an embedded code signal by 0 along with the storage of that embedded code contains a kind of inefficiency.

It is preferred to maintain the formalism of the 0 and 1 nature of the N -bit identification word, but
30 to have the 0's of the word induce a subtraction of their corresponding embedded code signal. Thus, in step 8 of Fig. 2, rather than only 'adding' the individual embedded code signals which correspond to a '1' in the N -bit identification word, we will also 'subtract' the individual embedded code signals which correspond to a '0' in the N -bit identification word.

- 16 -

At first glance this seems to add more apparent noise to the final composite signal. But it also increases the energy-wise separation of the 0's from the 1's, and thus the 'gain' which is applied in step 10, Fig. 2 can be correspondingly lower.

We can refer to this improvement as the use of true polarity. The main advantage of this improvement can largely be summarized as 'informational efficiency.'

'Perceptual Orthogonality' of the Individual Embedded Code Signals

The foregoing discussion contemplates the use of generally random noise-like signals as the individual embedded code signals. This is perhaps the simplest form of signal to generate. However, there is a form of informational optimization which can be applied to the set of the individual embedded signals, which the applicant describes under the rubric 'perceptual orthogonality.' This term is loosely based on the mathematical concept of the orthogonality of vectors, with the current additional requirement that this orthogonality should maximize the signal energy of the identification information while maintaining it below some perceptibility threshold. Put another way, the embedded code signals need not necessarily be random in nature.

Use and Improvements of the First Embodiment in the Field of Emulsion-Based Photography

The foregoing discussion outlined techniques that are applicable to photographic materials. The following section explores the details of this area further and discloses certain improvements which lend themselves to a broad range of applications.

The first area to be discussed involves the pre-application or pre-exposing of a serial number onto traditional photographic products, such as negative film, print paper, transparencies, etc. In general, this is a way to embed *a priori* unique serial numbers (and by implication, ownership and tracking information) into photographic material. The serial numbers themselves would be a permanent part of the normally exposed picture, as opposed to being relegated to the margins or stamped on the back of a printed photograph, which all require separate locations and separate methods of copying. The 'serial number' as it is called here is generally synonymous with the N-bit identification word, only now we are using a more common industrial terminology.

In Fig. 2, step 11, the disclosure calls for the storage of the "original [image]" along with code images. Then in Fig. 3, step 9, it directs that the original be subtracted from the suspect image, thereby leaving the possible identification codes plus whatever noise and corruption has accumulated. Therefore, the previous disclosure made the tacit assumption that there exists an original without the composite embedded signals.

Now in the case of selling print paper and other duplication film products, this will still be the case, i.e., an "original" without the embedded codes will indeed exist and the basic methodology of the first embodiment can be employed. The original film serves perfectly well as an 'unencoded original.'

However, in the case where pre-exposed negative film is used, the composite embedded signal
5 pre-exists on the original film and thus there will never be an "original" separate from the pre-embedded signal. It is this latter case, therefore, which will be examined a bit more closely, along with observations on how to best use the principles discussed above (the former cases adhering to the previously outlined methods).

The clearest point of departure for the case of pre-numbered negative film, i.e. negative film
10 which has had each and every frame pre-exposed with a very faint and unique composite embedded signal, comes at step 9 of Fig. 3 as previously noted. There are certainly other differences as well, but they are mostly logistical in nature, such as how and when to embed the signals on the film, how to store the code numbers and serial number, etc. Obviously the pre-exposing of film would involve a major change to the general mass production process of creating and packaging film.

Fig. 4 has a schematic outlining one potential post-hoc mechanism for pre-exposing film. 'Post-hoc' refers to applying a process after the full common manufacturing process of film has already taken place. Eventually, economies of scale may dictate placing this pre-exposing process directly into the chain of manufacturing film. Depicted in Fig. 4 is what is commonly known as a film writing system. The computer, 106, displays the composite signal produced in step 8, Fig. 2, on its phosphor screen. A given
20 frame of film is then exposed by imaging this phosphor screen, where the exposure level is generally very faint, i.e. generally imperceptible. Clearly, the marketplace will set its own demands on how faint this should be, that is, the level of added 'graininess' as practitioners would put it. Each frame of film is sequentially exposed, where in general the composite image displayed on the CRT 102 is changed for each and every frame, thereby giving each frame of film a different serial number. The transfer lens 104
25 highlights the focal conjugate planes of a film frame and the CRT face.

Getting back to the applying the principles of the foregoing embodiment in the case of pre-exposed negative film... At step 9, Fig. 3, if we were to subtract the "original" with its embedded code, we would obviously be "erasing" the code as well since the code is an integral part of the original. Fortunately, remedies do exist and identifications can still be made. However, it will be a challenge to artisans who
30 refine this embodiment to have the signal to noise ratio of the identification process in the pre-exposed negative case approach the signal to noise ratio of the case where the un-encoded original exists.

A succinct definition of the problem is in order at this point. Given a suspect picture (signal), find the embedded identification code IF a code exists at all. The problem reduces to one of finding the amplitude of each and every individual embedded code signal within the suspect picture, not only within
35 the context of noise and corruption as was previously explained, but now also within the context of the

- 18 -

coupling between a captured image and the codes. 'Coupling' here refers to the idea that the captured image "randomly biases" the cross-correlation.

So, bearing in mind this additional item of signal coupling, the identification process now estimates the signal amplitude of each and every individual embedded code signal (as opposed to taking the cross-correlation result of step 12, Fig. 3). If our identification signal exists in the suspect picture, the amplitudes thus found will split into a polarity with positive amplitudes being assigned a '1' and negative amplitudes being assigned a '0'. Our unique identification code manifests itself. If, on the other hand, no such identification code exists or it is someone else's code, then a random gaussian-like distribution of amplitudes is found with a random hash of values.

It remains to provide a few more details on how the amplitudes of the individual embedded codes are found. Again, fortunately, this exact problem has been treated in other technological applications. Besides, throw this problem and a little food into a crowded room of mathematicians and statisticians and surely a half dozen optimized methodologies will pop out after some reasonable period of time. It is a rather cleanly defined problem.

One specific example solution comes from the field of astronomical imaging. Here, it is a mature prior art to subtract out a "thermal noise frame" from a given CCD image of an object. Often, however, it is not precisely known what scaling factor to use in subtracting the thermal frame, and a search for the correct scaling factor is performed. This is precisely the task of this step of the present embodiment.

General practice merely performs a common search algorithm on the scaling factor, where a scaling factor is chosen and a new image is created according to:

$$\text{NEW IMAGE} = \text{ACQUIRED IMAGE} - \text{SCALE} * \text{THERMAL IMAGE} \quad (4)$$

The new image is applied to the fast fourier transform routine and a scale factor is eventually found which minimizes the integrated high frequency content of the new image. This general type of search operation with its minimization of a particular quantity is exceedingly common. The scale factor thus found is the sought-for "amplitude." Refinements which are contemplated but not yet implemented are where the coupling of the higher derivatives of the acquired image and the embedded codes are estimated and removed from the calculated scale factor. In other words, certain bias effects from the coupling mentioned earlier are present and should be eventually accounted for and removed both through theoretical and empirical experimentation.

Use and Improvements in the Detection of Signal or Image Alteration

Apart from the basic need of identifying a signal or image as a whole, there is also a rather ubiquitous need to detect possible alterations to a signal or image. The following section describes how the foregoing embodiment, with certain modifications and improvements, can be used as a powerful tool in this area. The potential scenarios and applications of detecting alterations are innumerable.

- 19 -

To first summarize, assume that we have a given signal or image which has been positively identified using the basic methods outlined above. In other words, we know its N-bit identification word, its individual embedded code signals, and its composite embedded code. We can then fairly simply create a spatial map of the composite code's amplitude within our given signal or image. Furthermore, we can
5 divide this amplitude map by the known composite code's spatial amplitude, giving a normalized map, i.e. a map which should fluctuate about some global mean value. By simple examination of this map, we can visually detect any areas which have been significantly altered wherein the value of the normalized amplitude dips below some statistically set threshold based purely on typical noise and corruption (error).

The details of implementing the creation of the amplitude map have a variety of choices. One is
10 to perform the same procedure which is used to determine the signal amplitude as described above, only now we step and repeat the multiplication of any given area of the signal/image with a gaussian weight function centered about the area we are investigating.

Universal Versus Custom Codes

15 The disclosure thus far has outlined how each and every source signal has its own unique set of individual embedded code signals. This entails the storage of a significant amount of additional code information above and beyond the original, and many applications may merit some form of economizing.

One such approach to economizing is to have a given set of individual embedded code signals be common to a batch of source materials. For example, one thousand images can all utilize the same basic
20 set of individual embedded code signals. The storage requirements of these codes then become a small fraction of the overall storage requirements of the source material.

Furthermore, some applications can utilize a universal set of individual embedded code signals, i.e., codes which remain the same for all instances of distributed material. This type of requirement would be seen by systems which wish to hide the N-bit identification word itself, yet have standardized equipment
25 be able to read that word. This can be used in systems which make go/no go decisions at point-of-read locations. The potential drawback to this set-up is that the universal codes are more prone to be sleuthed or stolen; therefore they will not be as secure as the apparatus and methodology of the previously disclosed arrangement. Perhaps this is just the difference between 'high security' and 'air-tight security,' a distinction carrying little weight with the bulk of potential applications.

30

Use in Printing, Paper, Documents, Plastic Coated Identification Cards, and Other Material Where Global Embedded Codes Can Be Imprinted

The term 'signal' is often used narrowly to refer to digital data information, audio signals, images, etc. A broader interpretation of 'signal,' and the one more generally intended, includes any form
35 of modulation of any material whatsoever. Thus, the micro-topology of a piece of common paper becomes

- 20 -

a 'signal' (e.g. its height as a function of x-y coordinates). The reflective properties of a flat piece of plastic (as a function of space also) becomes a signal. The point is that photographic emulsions, audio signals, and digitized information are not the only types of signals capable of utilizing the principles described herein.

5 As a case in point, a machine very much resembling a braille printing machine can be designed so as to imprint unique 'noise-like' indentations as outlined above. These indentations can be applied with a pressure which is much smaller than is typically applied in creating braille, to the point where the patterns are not noticed by a normal user of the paper. But by following the steps of the present disclosure and applying them via the mechanism of micro-indentations, a unique identification code can be placed onto
10 any given sheet of paper, be it intended for everyday stationary purposes, or be it for important documents, legal tender, or other secured material.

 The reading of the identification material in such an embodiment generally proceeds by merely reading the document optically at a variety of angles. This would become an inexpensive method for deducing the micro-topology of the paper surface. Certainly other forms of reading the topology of the
15 paper are possible as well.

 In the case of plastic encased material such as identification cards, e.g. driver's licenses, a similar braille-like impressions machine can be utilized to imprint unique identification codes. Subtle layers of photoreactive materials can also be embedded inside the plastic and 'exposed.'

 It is clear that wherever a material exists which is capable of being modulated by 'noise-like'
20 signals, that material is an appropriate carrier for unique identification codes and utilization of the principles disclosed herein. All that remains is the matter of economically applying the identification information and maintaining the signal level below an acceptability threshold which each and every application will define for itself.

25 REAL TIME ENCODER

 While the first class of embodiments most commonly employs a standard microprocessor or computer to perform the encodation of an image or signal, it is possible to utilize a custom encodation device which may be faster than a typical Von Neuman-type processor. Such a system can be utilized with all manner of serial data streams.

30 Music and videotape recordings are examples of serial data streams -- data streams which are often pirated. It would assist enforcement efforts if authorized recordings were encoded with identification data so that pirated knock-offs could be traced to the original from which they were made.

 Piracy is but one concern driving the need for applicant's technology. Another is authentication. Often it is important to confirm that a given set of data is really what it is purported to be (often several
35 years after its generation).

- 21 -

To address these and other needs, the system 200 of Fig. 5 can be employed. System 200 can be thought of as an identification coding black box 202. The system 200 receives an input signal (sometimes termed the "master" or "unencoded" signal) and a code word, and produces (generally in real time) an identification-coded output signal. (Usually, the system provides key data for use in later decoding.)

5 The contents of the "black box" 202 can take various forms. An exemplary black box system is shown in Fig. 6 and includes a look-up table 204, a digital noise source 206, first and second scalars 208, 210, an adder/subtractor 212, a memory 214, and a register 216.

10 The input signal (which in the illustrated embodiment is an 8 - 20 bit data signal provided at a rate of one million samples per second, but which in other embodiments could be an analog signal if appropriate A/D and D/A conversion is provided) is applied from an input 218 to the address input 220 of the look-up table 204. For each input sample (i.e. look-up table address), the table provides a corresponding 8-bit digital output word. This output word is used as a scaling factor that is applied to one input of the first scalar 208.

15 The first scalar 208 has a second input, to which is applied an 8-bit digital noise signal from source 206. (In the illustrated embodiment, the noise source 206 comprises an analog noise source 222 and an analog-to-digital converter 224 although, again, other implementations can be used.) The noise source in the illustrated embodiment has a zero mean output value, with a full width half maximum (FWHM) of 50 - 100 digital numbers (e.g. from -75 to +75).

20 The first scalar 208 multiplies the two 8-bit words at its inputs (scale factor and noise) to produce -- for each sample of the system input signal -- a 16-bit output word. Since the noise signal has a zero mean value, the output of the first scalar likewise has a zero mean value.

25 The output of the first scalar 208 is applied to the input of the second scalar 210. The second scalar serves a global scaling function, establishing the absolute magnitude of the identification signal that will ultimately be embedded into the input data signal. The scaling factor is set through a scale control device 226 (which may take a number of forms, from a simple rheostat to a graphically implemented control in a graphical user interface), permitting this factor to be changed in accordance with the requirements of different applications. The second scalar 210 provides on its output line 228 a scaled noise signal. Each sample of this scaled noise signal is successively stored in the memory 214.

30 (In the illustrated embodiment, the output from the first scalar 208 may range between -1500 and +1500 (decimal), while the output from the second scalar 210 is in the low single digits, (such as between -2 and +2).)

 Register 216 stores a multi-bit identification code word. In the illustrated embodiment this code word consists of 8 bits, although larger code words (up to hundreds of bits) are commonly used. These bits are referenced, one at a time, to control how the input signal is modulated with the scaled noise signal.

- 22 -

In particular, a pointer 230 is cycled sequentially through the bit positions of the code word in register 216 to provide a control bit of "0" or "1" to a control input 232 of the adder/subtractor 212. If, for a particular input signal sample, the control bit is a "1", the scaled noise signal sample on line 232 is added to the input signal sample. If the control bit is a "0", the scaled noise signal sample is subtracted from the input signal sample. The output 234 from the adder/subtractor 212 provides the black box's output signal.

The addition or subtraction of the scaled noise signal in accordance with the bits of the code word effects a modulation of the input signal that is generally imperceptible. However, with knowledge of the contents of the memory 214, a user can later decode the encoding, determining the code number used in the original encoding process. (Actually, use of memory 214 is optional, as explained below.)

It will be recognized that the encoded signal can be distributed in well known ways, including converted to printed image form, stored on magnetic media (floppy diskette, analog or DAT tape, etc.), CD-ROM, etc. etc.

15 Decoding

A variety of techniques can be used to determine the identification code with which a suspect signal has been encoded. Two are discussed below. The first is less preferable than the latter for most applications, but is discussed herein so that the reader may have a fuller context within which to understand the disclosed technology.

20 More particularly, the first decoding method is a difference method, relying on subtraction of corresponding samples of the original signal from the suspect signal to obtain difference samples, which are then examined (typically individually) for deterministic coding indicia (i.e. the stored noise data). This approach may thus be termed a "sample-based, deterministic" decoding technique.

The second decoding method does not make use of the original signal. Nor does it examine 25 particular samples looking for predetermined noise characteristics. Rather, the statistics of the suspect signal (or a portion thereof) are considered in the aggregate and analyzed to discern the presence of identification coding that permeates the entire signal. The reference to permeation means the entire identification code can be discerned from a small fragment of the suspect signal. This latter approach may thus be termed a "holographic, statistical" decoding technique.

30 Both of these methods begin by registering the suspect signal to match the original. This entails scaling (e.g. in amplitude, duration, color balance, etc.), and sampling (or resampling) to restore the original sample rate. As in the earlier described embodiment, there are a variety of well understood techniques by which the operations associated with this registration function can be performed.

As noted, the first decoding approach proceeds by subtracting the original signal from the 35 registered, suspect signal, leaving a difference signal. The polarity of successive difference signal samples

- 23 -

can then be compared with the polarities of the corresponding stored noise signal samples to determine the identification code. That is, if the polarity of the first difference signal sample matches that of the first noise signal sample, then the first bit of the identification code is a "1." (In such case, the polarity of the 9th, 17th, 25th, etc. samples should also all be positive.) If the polarity of the first difference signal sample is opposite that of the corresponding noise signal sample, then the first bit of the identification code is a "0."

By conducting the foregoing analysis with eight successive samples of the difference signal, the sequence of bits that comprise the original code word can be determined. If, as in the illustrated embodiment, pointer 230 stepped through the code word one bit at a time, beginning with the first bit, during encoding, then the first 8 samples of the difference signal can be analyzed to uniquely determine the value of the 8-bit code word.

In a noise-free world (speaking here of noise independent of that with which the identification coding is effected), the foregoing analysis would always yield the correct identification code. But a process that is only applicable in a noise-free world is of limited utility indeed.

(Further, accurate identification of signals in noise-free contexts can be handled in a variety of other, simpler ways: e.g. checksums; statistically improbable correspondence between suspect and original signals; etc.)

While noise-induced aberrations in decoding can be dealt with -- to some degree -- by analyzing large portions of the signal, such aberrations still place a practical ceiling on the confidence of the process. Further, the villain that must be confronted is not always as benign as random noise. Rather, it increasingly takes the form of human-caused corruption, distortion, manipulation, etc. In such cases, the desired degree of identification confidence can only be achieved by other approaches.

The illustrated embodiment (the "holographic, statistical" decoding technique) relies on recombining the suspect signal with certain noise data (typically the data stored in memory 214), and analyzing the entropy of the resulting signal. "Entropy" need not be understood in its most strict mathematical definition, it being merely the most concise word to describe randomness (noise, smoothness, snowiness, etc.).

Most serial data signals are not random. That is, one sample usually correlates -- to some degree -- with the adjacent samples. Noise, in contrast, typically is random. If a random signal (e.g. noise) is added to (or subtracted from) a non-random signal, the entropy of the resulting signal generally increases. That is, the resulting signal has more random variations than the original signal. This is the case with the encoded output signal produced by the present encoding process; it has more entropy than the original, unencoded signal.

If, in contrast, the addition of a random signal to (or subtraction from) a non-random signal reduces entropy, then something unusual is happening. It is this anomaly that the present decoding process uses to detect embedded identification coding.

To fully understand this entropy-based decoding method, it is first helpful to highlight a
5 characteristic of the original encoding process: the similar treatment of every eighth sample.

In the encoding process discussed above, the pointer 230 increments through the code word, one bit for each successive sample of the input signal. If the code word is eight bits in length, then the pointer returns to the same bit position in the code word every eighth signal sample. If this bit is a "1", noise is added to the input signal; if this bit is a "0", noise is subtracted from the input signal. Due to the cyclic
10 progression of the pointer 230, every eighth sample of an encoded signal thus shares a characteristic: they are all either augmented by the corresponding noise data (which may be negative), or they are all diminished, depending on whether the bit of the code word then being addressed by pointer 230 is a "1" or a "0".

To exploit this characteristic, the entropy-based decoding process treats every eighth sample of the
15 suspect signal in like fashion. In particular, the process begins by adding to the 1st, 9th, 17th, 25th, etc. samples of the suspect signal the corresponding scaled noise signal values stored in the memory 214 (i.e. those stored in the 1st, 9th, 17th, 25th, etc., memory locations, respectively). The entropy of the resulting signal (i.e. the suspect signal with every 8th sample modified) is then computed.

(Computation of a signal's entropy or randomness is well understood by artisans in this field.
20 One generally accepted technique is to take the derivative of the signal at each sample point, square these values, and then sum over the entire signal. However, a variety of other well known techniques can alternatively be used.)

The foregoing step is then repeated, this time subtracting the stored noise values from the 1st, 9th, 17th, 25 etc. suspect signal samples.

25 One of these two operations will undo the encoding process and reduce the resulting signal's entropy; the other will aggravate it. If adding the noise data in memory 214 to the suspect signal reduces its entropy, then this data must earlier have been subtracted from the original signal. This indicates that pointer 230 was pointing to a "0" bit when these samples were encoded. (A "0" at the control input of adder/subtractor 212 caused it to subtract the scaled noise from the input signal.)

30 Conversely, if subtracting the noise data from every eighth sample of the suspect signal reduces its entropy, then the encoding process must have earlier added this noise. This indicates that pointer 230 was pointing to a "1" bit when samples 1, 9, 17, 25, etc., were encoded.

By noting whether entropy decreases by (a) adding or (b) subtracting the stored noise data to/from the suspect signal, it can be determined that the first bit of the code word is (a) a "0", or (b) a "1".

- 25 -

The foregoing operations are then conducted for the group of spaced samples of the suspect signal beginning with the second sample (i.e. 2, 10, 18, 26 ...). The entropy of the resulting signals indicate whether the second bit of the code word is a "0" or a "1". Likewise with the following 6 groups of spaced samples in the suspect signal, until all 8 bits of the code word have been discerned.

5 It will be appreciated that the foregoing approach is not sensitive to corruption mechanisms that alter the values of individual samples; instead, the process considers the entropy of the signal as a whole, yielding a high degree of confidence in the results. Further, even small excerpts of the signal can be analyzed in this manner, permitting piracy of even small details of an original work to be detected. The results are thus statistically robust, both in the face of natural and human corruption of the suspect signal.

10 It will further be appreciated that the use of an N-bit code word in this real time embodiment provides benefits analogous to those discussed above in connection with the batch encoding system. (Indeed, the present embodiment may be conceptualized as making use of N different noise signals, just as in the batch encoding system. The first noise signal is a signal having the same extent as the input signal, and comprising the scaled noise signal at the 1st, 9th, 17th, 25th, etc., samples (assuming N=8), with
15 zeroes at the intervening samples. The second noise signal is a similar one comprising the scaled noise signal at the 2d, 10th, 18th, 26th, etc., samples, with zeroes at the intervening samples. Etc. These signals are all combined to provide a composite noise signal.) One of the important advantages inherent in such a system is the high degree of statistical confidence (confidence which doubles with each successive bit of the identification code) that a match is really a match. The system does not rely on subjective
20 evaluation of a suspect signal for a single, deterministic embedded code signal.

Illustrative Variations

From the foregoing description, it will be recognized that numerous modifications can be made to the illustrated systems without changing the fundamental principles. A few of these variations are
25 described below.

The above-described decoding process tries both adding and subtracting stored noise data to/from the suspect signal in order to find which operation reduces entropy. In other embodiments, only one of these operations needs to be conducted. For example, in one alternative decoding process the stored noise data corresponding to every eighth sample of the suspect signal is only added to said samples. If the
30 entropy of the resulting signal is thereby increased, then the corresponding bit of the code word is a "1" (i.e. this noise was added earlier, during the encoding process, so adding it again only compounds the signal's randomness). If the entropy of the resulting signal is thereby decreased, then the corresponding bit of the code word is a "0". A further test of entropy if the stored noise samples are subtracted is not required.

- 26 -

The statistical reliability of the identification process (coding and decoding) can be designed to exceed virtually any confidence threshold (e.g. 99.9%, 99.99%, 99.999%, etc. confidence) by appropriate selection of the global scaling factors, etc. Additional confidence in any given application (unnecessary in most applications) can be achieved by rechecking the decoding process.

5 One way to recheck the decoding process is to remove the stored noise data from the suspect signal in accordance with the bits of the discerned code word, yielding a "restored" signal (e.g. if the first bit of the code word is found to be "1," then the noise samples stored in the 1st, 9th, 17th, etc. locations of the memory 214 are subtracted from the corresponding samples of the suspect signal). The entropy of the restored signal is measured and used as a baseline in further measurements. Next, the process is
10 repeated, this time removing the stored noise data from the suspect signal in accordance with a modified code word. The modified code word is the same as the discerned code word, except 1 bit is toggled (e.g. the first). The entropy of the resulting signal is determined, and compared with the baseline. If the toggling of the bit in the discerned code word resulted in increased entropy, then the accuracy of that bit of the discerned code word is confirmed. The process repeats, each time with a different bit of the discerned
15 code word toggled, until all bits of the code word have been so checked. Each change should result in an increase in entropy compared to the baseline value.

The data stored in memory 214 is subject to a variety of alternatives. In the foregoing discussion, memory 214 contains the scaled noise data. In other embodiments, the unscaled noise data can be stored instead.

20 In still other embodiments, it can be desirable to store at least part of the input signal itself in memory 214. For example, the memory can allocate 8 signed bits to the noise sample, and 16 bits to store the most significant bits of an 18- or 20-bit audio signal sample. This has several benefits. One is that it simplifies registration of a "suspect" signal. Another is that, in the case of encoding an input signal which was already encoded, the data in memory 214 can be used to discern which of the encoding processes was
25 performed first. That is, from the input signal data in memory 214 (albeit incomplete), it is generally possible to determine with which of two code words it has been encoded.

Yet another alternative for memory 214 is that it can be omitted altogether.

One way this can be achieved is to use a deterministic noise source in the encoding process, such as an algorithmic noise generator seeded with a known key number. The same deterministic noise source,
30 seeded with the same key number, can be used in the decoding process. In such an arrangement, only the key number needs be stored for later use in decoding, instead of the large data set usually stored in memory 214.

Alternatively, if the noise signal added during encoding does not have a zero mean value, and the length N of the code word is known to the decoder, then a universal decoding process can be implemented.
35 This process uses the same entropy test as the foregoing procedures, but cycles through possible code

words, adding/subtracting a small dummy noise value (e.g. less than the expected mean noise value) to every Nth sample of the suspect signal, in accordance with the bits of the code word being tested, until a reduction in entropy is noted. Such an approach is not favored for most applications, however, because it offers less security than the other embodiments (e.g. it is subject to cracking by brute force).

5 Many applications are well served by the embodiment illustrated in Fig. 7, in which different code words are used to produce several differently encoded versions of an input signal, each making use of the same noise data. More particularly, the embodiment 240 of Fig. 7 includes a noise store 242 into which noise from source 206 is written during the identification-coding of the input signal with a first code word. (The noise source of Fig. 7 is shown outside of the real time encoder 202 for convenience of illustration.)

10 Thereafter, additional identification-coded versions of the input signal can be produced by reading the stored noise data from the store and using it in conjunction with second through Nth code words to encode the signal. (While binary-sequential code words are illustrated in Fig. 7, in other embodiments arbitrary sequences of code words can be employed.) With such an arrangement, a great number of differently-encoded signals can be produced, without requiring a proportionally-sized long term noise memory.

15 Instead, a fixed amount of noise data is stored, whether encoding an original once or a thousand times.

(If desired, several differently-coded output signals can be produced at the same time, rather than serially. One such implementation includes a plurality of adder/subtractor circuits 212, each driven with the same input signal and with the same scaled noise signal, but with different code words. Each, then, produces a differently encoded output signal.)

20 In applications having a great number of differently-encoded versions of the same original, it will be recognized that the decoding process need not always discern every bit of the code word. Sometimes, for example, the application may require identifying only a group of codes to which the suspect signal belongs. (E.g., high order bits of the code word might indicate an organization to which several differently coded versions of the same source material were provided, with low-order bits identifying specific copies. To identify the organization with which a suspect signal is associated, it may not be necessary to examine the low order bits, since the organization can be identified by the high order bits alone.) If the identification requirements can be met by discerning a subset of the code word bits in the suspect signal, the decoding process can be shortened.

25 Some applications may be best served by restarting the encoding process -- sometimes with a different code word -- several times within an integral work. Consider, as an example, videotaped productions (e.g. television programming). Each frame of a videotaped production can be identification-coded with a unique code number, processed in real-time with an arrangement 248 like that shown in Fig. 8. Each time a vertical retrace is detected by sync detector 250, the noise source 206 resets (e.g. to repeat the sequence just produced) and an identification code increments to the next value. Each frame of the

videotape is thereby uniquely identification-coded. Typically, the encoded signal is stored on a videotape for long term storage (although other storage media, including laser disks, can be used).

Returning to the encoding apparatus, the look-up table 204 in the illustrated embodiment exploits the fact that high amplitude samples of the input data signal can tolerate (without objectionable degradation of the output signal) a higher level of encoded identification coding than can low amplitude input samples. Thus, for example, input data samples having decimal values of 0, 1 or 2 may correspond (in the look-up table 204) to scale factors of unity (or even zero), whereas input data samples having values in excess of 200 may correspond to scale factors of 15. Generally speaking, the scale factors and the input sample values correspond by a square root relation. That is, a four-fold increase in a value of the sampled input signal corresponds to approximately a two-fold increase in a value of the scaling factor associated therewith.

(The parenthetical reference to zero as a scaling factor alludes to cases, e.g., in which the source signal is temporally or spatially devoid of information content. In an image, for example, a region characterized by several contiguous sample values of zero may correspond to a jet black region of the frame. A scaling value of zero may be appropriate here since there is essentially no image data to be pirated.)

Continuing with the encoding process, those skilled in the art will recognize the potential for "rail errors" in the illustrated embodiment. For example, if the input signal consists of 8-bit samples, and the samples span the entire range from 0 to 255 (decimal), then the addition or subtraction of scaled noise to/from the input signal may produce output signals that cannot be represented by 8 bits (e.g. -2, or 257). A number of well-understood techniques exist to rectify this situation, some of them proactive and some of them reactive. (Among these known techniques are: specifying that the input signal shall not have samples in the range of 0-4 or 251-255, thereby safely permitting modulation by the noise signal; or including provision for detecting and adaptively modifying input signal samples that would otherwise cause rail errors.)

While the illustrated embodiment describes stepping through the code word sequentially, one bit at a time, to control modulation of successive bits of the input signal, it will be appreciated that the bits of the code word can be used other than sequentially for this purpose. Indeed, bits of the code word can be selected in accordance with any predetermined algorithm.

The dynamic scaling of the noise signal based on the instantaneous value of the input signal is an optimization that can be omitted in many embodiments. That is, the look-up table 204 and the first scaler 208 can be omitted entirely, and the signal from the digital noise source 206 applied directly (or through the second, global scaler 210) to the adder/subtractor 212.

It will be further recognized that the use of a zero-mean noise source simplifies the illustrated embodiment, but is not essential. A noise signal with another mean value can readily be used, and D.C. compensation (if needed) can be effected elsewhere in the system.

The use of a noise source 206 is also optional. A variety of other signal sources can be used, depending on application- dependent constraints (e.g. the threshold at which the encoded identification signal becomes perceptible). In many instances, the level of the embedded identification signal is low enough that the identification signal needn't have a random aspect; it is imperceptible regardless of its nature. A pseudo random source 206, however, is usually desired because it provides the greatest identification code signal S/N ratio (a somewhat awkward term in this instance) for a level of imperceptibility of the embedded identification signal.

It will be recognized that identification coding need not occur after a signal has been reduced to stored form as data (i.e. "fixed in tangible form," in the words of the U.S. Copyright Act). Consider, for example, the case of popular musicians whose performances are often recorded illicitly. By identification coding the audio before it drives concert hall speakers, unauthorized recordings of the concert can be traced to a particular place and time. Likewise, live audio sources such as 911 emergency calls can be encoded prior to recording so as to facilitate their later authentication.

While the black box embodiment has been described as a stand alone unit, it will be recognized that it can be integrated into a number of different tools/instruments as a component. One is a scanner, which can embed identification codes in the scanned output data. (The codes can simply serve to memorialize that the data was generated by a particular scanner). Another is in creativity software, such as popular drawing/graphics/animation/paint programs offered by Adobe, Macromedia, Corel, and the like.

Finally, while the real-time encoder 202 has been illustrated with reference to a particular hardware implementation, it will be recognized that a variety of other implementations can alternatively be employed. Some utilize other hardware configurations. Others make use of software routines for some or all of the illustrated functional blocks. (The software routines can be executed on any number of different general purpose programmable computers, such as 80x86 PC-compatible computers, RISC-based workstations, etc.)

TYPES OF NOISE, QUASI-NOISE, AND OPTIMIZED-NOISE

Heretofore this disclosure posulated Gaussian noise, "white noise," and noise generated directly from application instrumentation as a few of the many examples of the kind of carrier signal appropriate to carry a single bit of information throughout an image or signal. It is possible to be even more proactive in "designing" characteristics of noise in order to achieve certain goals. The "design" of using Gaussian or instrumental noise was aimed somewhat toward "absolute" security. This section of the disclosure takes a

- 30 -

look at other considerations for the design of the noise signals which may be considered the ultimate carriers of the identification information.

For some applications it might be advantageous to design the noise carrier signal (e.g. the Nth embedded code signal in the first embodiment; the scaled noise data in the second embodiment), so as to provide more absolute signal strength to the identification signal relative to the perceptibility of that signal. One example is the following. It is recognized that a true Gaussian noise signal has the value '0' occur most frequently, followed by 1 and -1 at equal probabilities to each other but lower than '0', 2 and -2 next, and so on. Clearly, the value zero carries no information as it is used in such an embodiment. Thus, one simple adjustment, or design, would be that any time a zero occurs in the generation of the embedded code signal, a new process takes over, whereby the value is converted "randomly" to either a 1 or a -1. In logical terms, a decision would be made: if '0', then random(1,-1). The histogram of such a process would appear as a Gaussian/Poissonian type distribution, except that the 0 bin would be empty and the 1 and -1 bin would be increased by half the usual histogram value of the 0 bin.

In this case, identification signal energy would always be applied at all parts of the signal. A few of the trade-offs include: there is a (probably negligible) lowering of security of the codes in that a "deterministic component" is a part of generating the noise signal. The reason this might be completely negligible is that we still wind up with a coin flip type situation on randomly choosing the 1 or the -1. Another trade-off is that this type of designed noise will have a higher threshold of perceptibility, and will only be applicable to applications where the least significant bit of a data stream or image is already negligible relative to the commercial value of the material, i.e. if the least significant bit were stripped from the signal (for all signal samples), no one would know the difference and the value of the material would not suffer. This blocking of the zero value in the example above is but one of many ways to "optimize" the noise properties of the signal carrier, as anyone in the art can realize. We refer to this also as "quasi-noise" in the sense that natural noise can be transformed in a pre-determined way into signals which for all intents and purposes will read as noise. Also, cryptographic methods and algorithms can easily, and often by definition, create signals which are perceived as completely random. Thus the word "noise" can have different connotations, primarily between that as defined subjectively by an observer or listener, and that defined mathematically. The difference of the latter is that mathematical noise has different properties of security and the simplicity with which it can either be "sleuthed" or the simplicity with which instruments can "automatically recognize" the existence of this noise.

"Universal" Embedded Codes

The bulk of this disclosure teaches that for absolute security, the noise-like embedded code signals which carry the bits of information of the identification signal should be unique to each and every encoded signal, or, slightly less restrictive, that embedded code signals should be generated sparingly, such as using the

same embedded codes for a batch of 1000 pieces of film, for example. Be this as it may, there is a whole other approach to this issue wherein the use of what we will call "universal" embedded code signals can open up large new applications for this technology. The economics of these uses would be such that the de facto lowered security of these universal codes (e.g. they would be analyzable by time honored
5 cryptographic decoding methods, and thus potentially thwarted or reversed) would be economically negligible relative to the economic gains that the intended uses would provide. Piracy and illegitimate uses would become merely a predictable "cost" and a source of uncollected revenue only; a simple line item in an economic analysis of the whole. A good analogy of this is in the cable industry and the scrambling of video signals. Everybody seems to know that crafty, skilled technical individuals, who may be generally
10 law abiding citizens, can climb a ladder and flip a few wires in their cable junction box in order to get all the pay channels for free. The cable industry knows this and takes active measures to stop it and prosecute those caught, but the "lost revenue" derived from this practice remains prevalent but almost negligible as a percentage of profits gained from the scrambling system as a whole. The scrambling system as a whole is an economic success despite its lack of "absolute security."

15 The same holds true for applications of this technology wherein, for the price of lowering security by some amount, large economic opportunity presents itself. This section first describes what is meant by universal codes, then moves on to some of the interesting uses to which these codes can be applied.

Universal embedded codes generally refer to the idea that knowledge of the exact codes can be distributed. The embedded codes won't be put into a dark safe never to be touched until litigation arises
20 (as alluded to in other parts of this disclosure), but instead will be distributed to various locations where on-the-spot analysis can take place. Generally this distribution will still take place within a security controlled environment, meaning that steps will be taken to limit the knowledge of the codes to those with a need to know. Instrumentation which attempts to automatically detect copyrighted material is a non-human example of "something" with a need to know the codes.

25 There are many ways to implement the idea of universal codes, each with their own merits regarding any given application. For the purposes of teaching this art, we separate these approaches into three broad categories: universal codes based on libraries, universal codes based on deterministic formula, and universal codes based on pre-defined industry standard patterns. A rough rule of thumb is that the first is more secure than the latter two, but that the latter two are possibly more economical to implement
30 than the first.

Universal Codes: 1) Libraries of Universal Codes

The use of libraries of universal codes simply means that applicant's techniques are employed as described, except for the fact that only a limited set of the individual embedded code signals are generated
35 and that any given encoded material will make use of some sub-set of this limited "universal set." An

example is in order here. A photographic print paper manufacturer may wish to pre-expose every piece of 8 by 10 inch print paper which they sell with a unique identification code. They also wish to sell identification code recognition software to their large customers, service bureaus, stock agencies, and individual photographers, so that all these people can not only verify that their own material is correctly
5 ~~marked~~, but so that they can also determine if third party material which they are about to acquire has been identified by this technology as being copyrighted. This latter information will help them verify copyright holders and avoid litigation, among many other benefits. In order to "economically" institute this plan, they realize that generating unique individual embedded codes for each and every piece of print paper would generate Terabytes of independent information, which would need storing and to which
10 recognition software would need access. Instead, they decide to embed their print paper with 16 bit identification codes derived from a set of only 50 independent "universal" embedded code signals. The details of how this is done are in the next paragraph, but the point is that now their recognition software only needs to contain a limited set of embedded codes in their library of codes, typically on the order of 1 Megabyte to 10 Megabytes of information for 50x16 individual embedded codes splayed out onto an 8x10
15 photographic print (allowing for digital compression). The reason for picking 50 instead of just 16 is one of a little more added security, where if it were the same 16 embedded codes for all photographic sheets, not only would the serial number capability be limited to 2 to the 16th power, but lesser and lesser sophisticated pirates could crack the codes and remove them using software tools.

There are many different ways to implement this scheme, where the following is but one
20 exemplary method. It is determined by the wisdom of company management that a 300 pixels per inch criteria for the embedded code signals is sufficient resolution for most applications. This means that a composite embedded code image will contain 3000 pixels by 2400 pixels to be exposed at a very low level onto each 8x10 sheet. This gives 7.2 million pixels. Using our staggered coding system such as described in the black box implementation of Figs. 5 and 6, each individual embedded code signal will contain only
25 7.2 million divided by 16, or approximately 450K true information carrying pixels, i.e. every 16th pixel along a given raster line. These values will typically be in the range of 2 to -2 in digital numbers, or adequately described by a signed 3 bit number. The raw information content of an embedded code is then approximately 3/8th's bytes times 450K or about 170 Kilobytes. Digital compression can reduce this further. All of these decisions are subject to standard engineering optimization principles as defined by
30 any given application at hand, as is well known in the art. Thus we find that 50 of these independent embedded codes will amount to a few Megabytes. This is quite reasonable level to distribute as a "library" of universal codes within the recognition software. Advanced standard encryption devices could be employed to mask the exact nature of these codes if one were concerned that would-be pirates would buy the recognition software merely to reverse engineer the universal embedded codes. The recognition

- 33 -

software could simply unencrypt the codes prior to applying the recognition techniques taught in this disclosure.

The recognition software itself would certainly have a variety of features, but the core task it would perform is determining if there is some universal copyright code within a given image. The key questions become WHICH 16 of the total 50 universal codes it might contain, if any, and if there are 16 found, what are their bit values. The key variables in determining the answers to these questions are: registration, rotation, magnification (scale), and extent. In the most general case with no helpful hints whatsoever, all variables must be independently varied across all mutual combinations, and each of the 50 universal codes must then be checked by adding and subtracting to see if an entropy decrease occurs. Strictly speaking, this is an enormous job, but many helpful hints will be found which make the job much simpler, such as having an original image to compare to the suspected copy, or knowing the general orientation and extent of the image relative to an 8x10 print paper, which then through simple registration techniques can determine all of the variables to some acceptable degree. Then it merely requires cycling through the 50 universal codes to find any decrease in entropy. If one does, then 15 others should as well. A protocol needs to be set up whereby a given order of the 50 translates into a sequence of most significant bit through least significant bit of the ID code word. Thus if we find that universal code number "4" is present, and we find its bit value to be "0", and that universal codes "1" through "3" are definitely not present, then our most significant bit of our N-bit ID code number is a "0". Likewise, we find that the next lowest universal code present is number "7" and it turns out to be a "1", then our next most significant bit is a "1". Done properly, this system can cleanly trace back to the copyright owner so long as they registered their photographic paper stock serial number with some registry or with the manufacturer of the paper itself. That is, we look up in the registry that a paper using universal embedded codes 4,7,11,12,15,19,21,26,27,28,34,35,37,38,40, and 48, and having the embedded code 0110 0101 0111 0100 belongs to Leonardo de Boticelli, an unknown wildlife photographer and glacier cinematographer whose address is in Northern Canada. We know this because he dutifully registered his film and paper stock, a few minutes of work when he bought the stock, which he plopped into the "no postage necessary" envelope that the manufacturing company kindly provided to make the process ridiculously simple. Somebody owes Leonardo a royalty check it would appear, and certainly the registry has automated this royalty payment process as part of its services.

One final point is that truly sophisticated pirates and others with illicit intentions can indeed employ a variety of cryptographic and not so cryptographic methods to crack these universal codes, sell them, and make software and hardware tools which can assist in the removing or distorting of codes. We shall not teach these methods as part of this disclosure, however. In any event, this is one of the prices which must be paid for the ease of universal codes and the applications they open up.

Universal Codes: 2) Universal Codes Based on Deterministic Formulas

The libraries of universal codes require the storage and transmittal of Megabytes of independent, generally random data as the keys with which to unlock the existence and identity of signals and imagery that have been marked with universal codes. Alternatively, various deterministic formulas can be used which "generate" what appear to be random data/image frames, thereby obviating the need to store all of these codes in memory and interrogate each and of the "50" universal codes. Deterministic formulas can also assist in speeding up the process of determining the ID code once one is known to exist in a given signal or image. On the other hand, deterministic formulas lend themselves to sleuthing by less sophisticated pirates. And once sleuthed, they lend themselves to easier communication, such as posting on the Internet to a hundred newsgroups. There may well be many applications which do not care about sleuthing and publishing, and deterministic formulas for generating the individual universal embedded codes might be just the ticket.

Universal Codes: 3) "Simple" Universal Codes

This category is a bit of a hybrid of the first two, and is most directed at truly large scale implementations of the principles of this technology. The applications employing this class are of the type where staunch security is much less important than low cost, large scale implementation and the vastly larger economic benefits that this enables. One exemplary application is placement of identification recognition units directly within modestly priced home audio and video instrumentation (such as a TV). Such recognition units would typically monitor audio and/or video looking for these copyright identification codes, and thence triggering simple decisions based on the findings, such as disabling or enabling recording capabilities, or incrementing program specific billing meters which are transmitted back to a central audio/video service provider and placed onto monthly invoices. Likewise, it can be foreseen that "black boxes" in bars and other public places can monitor (listen with a microphone) for copyrighted materials and generate detailed reports, for use by ASCAP, BMI, and the like.

A core principle of simple universal codes is that some basic industry standard "noiselike" and seamlessly repetitive patterns are injected into signals, images, and image sequences so that inexpensive recognition units can either A) determine the mere existence of a copyright "flag", and B) additionally to A, determine precise identification information which can facilitate more complex decision making and actions.

In order to implement this particular embodiment, the basic principles of generating the individual embedded noise signals need to be simplified in order to accommodate inexpensive recognition signal processing circuitry, while maintaining the properties of effective randomness and holographic permeation. With large scale industry adoption of these simple codes, the codes themselves would border on public domain information (much as cable scrambling boxes are almost de facto public domain), leaving the door

- 35 -

open for determined pirates to develop black market countermeasures, but this situation would be quite analogous to the scrambling of cable video and the objective economic analysis of such illegal activity.

One prior art known to the applicant in this general area of pro-active copyright detection is the Serial Copy Management System adopted by many firms in the audio industry. To the best of applicant's knowledge, this system employs a non-audio "flag" signal which is not part of the audio data stream, but which is nevertheless grafted onto the audio stream and can indicate whether the associated audio data should or should not be duplicated. One problem with this system is that it is restricted to media and instrumentation which can support this extra "flag" signal. Another deficiency is that the flagging system carries no identity information which would be useful in making more complex decisions. Yet another difficulty is that high quality audio sampling of an analog signal can come arbitrarily close to making a perfect digital copy of some digital master and there seems to be no provision for inhibiting this possibility.

Applicant's technology can be brought to bear on these and other problems, in audio applications, video, and all of the other applications previously discussed. An exemplary application of simple universal codes is the following. A single industry standard "1.000000 second of noise" would be defined as the most basic indicator of the presence or absence of the copyright marking of any given audio signal. Fig. 9 has an example of what the waveform of an industry standard noise second might look like, both in the time domain 400 and the frequency domain 402. It is by definition a continuous function and would adapt to any combination of sampling rates and bit quantizations. It has a normalized amplitude and can be scaled arbitrarily to any digital signal amplitude. The signal level and the first M'th derivatives of the signal are continuous at the two boundaries 404 (Fig. 9C), such that when it is repeated, the "break" in the signal would not be visible (as a waveform) or audible when played through a high end audio system. The choice of 1 second is arbitrary in this example, where the precise length of the interval will be derived from considerations such as audibility, quasi-white noise status, seamless repeatability, simplicity of recognition processing, and speed with which a copyright marking determination can be made. The injection of this repeated noise signal onto a signal or image (again, at levels below human perception) would indicate the presence of copyright material. This is essentially a one bit identification code, and the embedding of further identification information will be discussed later on in this section. The use of this identification technique can extend far beyond the low cost home implementations discussed here, where studios could use the technique, and monitoring stations could be set up which literally monitor hundreds of channels of information simultaneously, searching for marked data streams, and furthermore searching for the associated identity codes which could be tied in with billing networks and royalty tracking systems.

This basic, standardized noise signature is seamlessly repeated over and over again and added to audio signals which are to be marked with the base copyright identification. Part of the reason for the word "simple" is seen here: clearly pirates will know about this industry standard signal, but their illicit

uses derived from this knowledge, such as erasure or corruption, will be economically minuscule relative to the economic value of the overall technique to the mass market. For most high end audio this signal will be some 80 to 100 dB down from full scale, or even much further; each situation can choose its own levels though certainly there will be recommendations. The amplitude of the signal can be modulated according to the audio signal levels to which the noise signature is being applied, i.e. the amplitude can increase significantly when a drum beats, but not so dramatically as to become audible or objectionable. These measures merely assist the recognition circuitry to be described.

Recognition of the presence of this noise signature by low cost instrumentation can be effected in a variety of ways. One rests on basic modifications to the simple principles of audio signal power metering. Software recognition programs can also be written, and more sophisticated mathematical detection algorithms can be applied to audio in order to make higher confidence detection identifications. In such embodiments, detection of the copyright noise signature involves comparing the time averaged power level of an audio signal with the time averaged power level of that same audio signal which has had the noise signature subtracted from it. If the audio signal with the noise signature subtracted has a lower power level than the unchanged audio signal, then the copyright signature is present and some status flag to that effect needs to be set. The main engineering subtleties involved in making this comparison include: dealing with audio speed playback discrepancies (e.g. an instrument might be 0.5% "slow" relative to exactly one second intervals); and, dealing with the unknown phase of the one second noise signature within any given audio (basically, its "phase" can be anywhere from 0 to 1 seconds). Another subtlety, not so central as the above two but which nonetheless should be addressed, is that the recognition circuits should not subtract a higher amplitude of the noise signature than was originally embedded onto the audio signal. Fortunately this can be accomplished by merely subtracting only a small amplitude of the noise signal, and if the power level goes down, this is an indication of "heading toward a trough" in the power levels. Yet another related subtlety is that the power level changes will be very small relative to the overall power levels, and calculations generally will need to be done with appropriate bit precision, e.g. 32 bit value operations and accumulations on 16-20 bit audio in the calculations of time averaged power levels.

Clearly, designing and packaging this power level comparison processing circuitry for low cost applications is an engineering optimization task. One trade-off will be the accuracy of making an identification relative to the "short-cuts" which can be made to the circuitry in order to lower its cost and complexity. One embodiment for placing this recognition circuitry inside of instrumentation is through a single programmable integrated circuit which is custom made for the task. Fig. 10 shows one such integrated circuit 506. Here the audio signal comes in, 500, either as a digital signal or as an analog signal to be digitized inside the IC 500, and the output is a flag 502 which is set to one level if the copyright noise signature is found, and to another level if it is not found. Also depicted is the fact that the

standardized noise signature waveform is stored in Read Only Memory, 504, inside the IC 506. There will be a slight time delay between the application of an audio signal to the IC 506 and the output of a valid flag 502, due to the need to monitor some finite portion of the audio before a recognition can place. In this case, there may need to be a "flag valid" output 508 where the IC informs the external world if it has had enough time to make a proper determination of the presence or absence of the copyright noise signature.

There are a wide variety of specific designs and philosophies of designs applied to accomplishing the basic function of the IC 506 of Fig. 10. Audio engineers and digital signal processing engineers are able to generate several fundamentally different designs. One such design is depicted in Fig. 11 by a process 599, which itself is subject to further engineering optimization as will be discussed. Fig. 11 depicts a flow chart for any of: an analog signal processing network, a digital signal processing network, or programming steps in a software program. We find an input signal 600 which along one path is applied to a time averaged power meter 602, and the resulting power output itself treated as a signal P_{avg} . To the upper right we find the standard noise signature 504 which will be read out at 125% of normal speed, 604, thus changing its pitch, giving the "pitch changed noise signal" 606. Then the input signal has this pitch changed noise signal subtracted in step 608, and this new signal is applied to the same form of time averaged power meter as in 602, here labelled 610. The output of this operation is also a time based signal here labelled as P_{pca} , 610. Step 612 then subtracts the power signal 602 from the power signal 610, giving an output difference signal P_{diff} , 613. If the universal standard noise signature does indeed exist on the input audio signal 600, then case 2, 616, will be created wherein a beat signal 618 of approximately 4 second period will show up on the output signal 613, and it remains to detect this beat signal with a step such as in Fig. 12, 622. Case 1, 614, is a steady noisy signal which exhibits no periodic beating. 125% at step 604 is chosen arbitrarily here, where engineering considerations would determine an optimal value, leading to different beat signal frequencies 618. Whereas waiting 4 seconds in this example would be quite a while, especially is you would want to detect at least two or three beats, Fig. 12 outlines how the basic design of Fig. 11 could be repeated and operated upon various delayed versions of the input signal, delayed by something like 1/20th of a second, with 20 parallel circuits working in concert each on a segment of the audio delayed by 0.05 seconds from their neighbors. In this way, a beat signal will show up approximately every 1/5th of a second and will look like a travelling wave down the columns of beat detection circuits. The existence or absence of this travelling beat wave triggers the detection flag 502. Meanwhile, there would be an audio signal monitor 624 which would ensure that, for example, at least two seconds of audio has been heard before setting the flag valid signal 508.

Though the audio example was described above, it should be clear to anyone in the art that the same type of definition of some repetitive universal noise signal or image could be applied to the many other signals, images, pictures, and physical media already discussed.

- 38 -

The above case deals only with a single bit plane of information, i.e., the noise signature signal is either there (1) or it isn't (0). For many applications, it would be nice to detect serial number information as well, which could then be used for more complex decisions, or for logging information on billing statements or whatnot. The same principles as the above would apply, but now there would be N independent noise signatures as depicted in Fig. 9 instead one single such signature. Typically, one such signature would be the master upon which the mere existence of a copyright marking is detected, and this would have generally higher power than the others, and then the other lower power "identification" noise signatures would be embedded into audio. Recognition circuits, once having found the existence of the primary noise signature, would then step through the other N noise signatures applying the same steps as described above. Where a beat signal is detected, this indicates the bit value of '1', and where no beat signal is detected, this indicates a bit value of '0'. It might be typical that N will equal 32, that way 2^{32} number of identification codes are available to any given industry employing this technology.

Use of this Technology When the Length of the Identification Code is 1

The principles described herein can obviously be applied in the case where only a single presence or absence of an identification signal – a fingerprint if you will – is used to provide confidence that some signal or image is copyrighted. The example above of the industry standard noise signature is one case in point. We no longer have the added confidence of the coin flip analogy, we no longer have tracking code capabilities or basic serial number capabilities, but many applications may not require these attributes and the added simplicity of a single fingerprint might outweigh these other attributes in any event.

The "Wallpaper" Analogy

The term "holographic" has been used in this disclosure to describe how an identification code number is distributed in a largely integral form throughout an encoded signal or image. This also refers to the idea that any given fragment of the signal or image contains the entire unique identification code number. As with physical implementations of holography, there are limitations on how small a fragment can become before one begins to lose this property, where the resolution limits of the holographic media are the main factor in this regard for holography itself. In the case of an uncorrupted distribution signal which has used the encoding device of Fig. 5, and which furthermore has used our "designed noise" of above wherein the zero's were randomly changed to a 1 or -1, then the extent of the fragment required is merely N contiguous samples in a signal or image raster line, where N is as defined previously being the length of our identification code number. This is an informational extreme; practical situations where noise and corruption are operative will require generally one, two or higher orders of magnitude more samples than this simple number N. Those skilled in the art will recognize that there are many variables

- 39 -

involved in pinning down precise statistics on the size of the smallest fragment with which an identification can be made.

For tutorial purposes, the applicant also uses the analogy that the unique identification code number is "wallpapered" across an image (or signal). That is, it is repeated over and over again all throughout an image. This repetition of the ID code number can be regular, as in the use of the encoder of Fig. 5, or random itself, where the bits in the ID code 216 of Fig. 6 are not stepped through in a normal repetitive fashion but rather are randomly selected on each sample, and the random selection stored along with the value of the output 228 itself. In any event, the information carrier of the ID code, the individual embedded code signal, does change across the image or signal. Thus as the wallpaper analogy summarizes: the ID code repeats itself over and over, but the patterns that each repetition imprints change randomly accordingly to a generally unsluthable key.

Lossy Data Compression

As earlier mentioned, applicant's preferred forms of identification coding withstand lossy data compression, and subsequent decompression. Such compression is finding increasing use, particularly in contexts such as the mass distribution of digitized entertainment programming (movies, etc.).

While data encoded according to the disclosed techniques can withstand all types of lossy compression known to applicant, those expected to be most commercially important are the CCITT G3, CCITT G4, JPEG, MPEG and JBIG compression/decompression standards. The CCITT standards are widely used in black-and-white document compression (e.g. facsimile and document-storage). JPEG is most widely used with still images. MPEG is most widely used with moving images. JBIG is a likely successor to the CCITT standards for use with black-and-white imagery. Such techniques are well known to those in the lossy data compression field; a good overview can be found in Pennebaker et al, *JPEG, Still Image Data Compression Standard*, Van Nostrand Reinhold, N.Y., 1993.

Towards Steganography Proper and the Use of this Technology in Passing More Complex Messages or Information

This disclosure concentrates on what above was called wallpapering a single identification code across an entire signal. This appears to be a desirable feature for many applications. However, there are other applications where it might be desirable to pass messages or to embed very long strings of pertinent identification information in signals and images. One of many such possible applications would be where a given signal or image is meant to be manipulated by several different groups, and that certain regions of an image are reserved for each group's identification and insertion of pertinent manipulation information.

In these cases, the code word 216 in Fig. 6 can actually change in some pre-defined manner as a function of signal or image position. For example, in an image, the code could change for each and every

- 40 -

raster line of the digital image. It might be a 16 bit code word, 216, but each scan line would have a new code word, and thus a 480 scan line image could pass a 980 (480 x 2 bytes) byte message. A receiver of the message would need to have access to either the noise signal stored in memory 214, or would have to know the universal code structure of the noise codes if that method of coding was being used. To the best of applicant's knowledge, this is a novel approach to the mature field of steganography.

In all three of the foregoing applications of universal codes, it will often be desirable to append a short (perhaps 8- or 16-bit) private code, which users would keep in their own secured places, in addition to the universal code. This affords the user a further modicum of security against potential erasure of the universal codes by sophisticated pirates.

Applicant's Prior Application

The Detailed Description to this point has simply repeated the disclosure of applicant's prior international application, laid open as PCT publication WO 95/14289. It was reproduced above simply to provide context for the disclosure which follows.

One Master Code Signal As A Distinction From N Independent Embedded Code Signals

In certain sections of this disclosure, perhaps exemplified in the section on the realtime encoder, an economizing step was taken whereby the N independent and source-signal-coextensive embedded code signals were so designed that the non-zero elements of any given embedded code signal were unique to just that embedded code signal and no others. Said more carefully, certain pixels/sample points of a given signal were "assigned" to some pre-determined m'th bit location in our N-bit identification word. Furthermore, and as another basic optimization of implementation, the aggregate of these assigned pixels/samples across all N embedded code signals is precisely the extent of the source signal, meaning each and every pixel/sample location in a source signal is assigned one and only one m'th bit place in our N-bit identification word. (This is not to say, however, that each and every pixel MUST be modified). As a matter of simplification we can then talk about a single master code signal (or "Snowy Image") rather than N independent signals, realizing that pre-defined locations in this master signal correspond to unique bit locations in our N-bit identification word. We therefore construct, via this circuitous route, this rather simple concept on the single master noise signal. Beyond mere economization and simplification, there are also performance reasons for this shift, primarily derived from the idea that individual bit places in our N-bit identification word are no longer "competing" for the information carrying capacity of a single pixel/sample.

With this single master more clearly understood, we can gain new insights into other sections of this disclosure and explore further details within the given application areas.

More of Deterministic Universal Codes Using the Master Code Concept

One case in point is to further explore the use of Deterministic Universal Codes, labelled as item "2" in the sections devoted to universal codes. A given user of this technology may opt for the following variant use of the principles of this technology. The user in question might be a mass distributor of home videos, but clearly the principles would extend to all other potential users of this technology. Fig. 13 pictorially represents the steps involved. In the example the user is one "Alien Productions." They first create an image canvas which is coextensive to the size of the video frames of their movie "Bud's Adventures." On this canvas they print the name of the movie, they place their logo and company name. Furthermore, they have specific information at the bottom, such as the distribution lot for the mass copying that they are currently cranking out, and as indicated, they actually have a unique frame number indicated. Thus we find the example of a standard image 700 which forms the initial basis for the creation of a master Snowy Image (master code signal) which will be added into the original movie frame, creating an output distributable frame. This image 700 can be either black & white or color. The process of turning this image 700 into a pseudo random master code signal is alluded to by the encryption/scrambling routine 702, wherein the original image 700 is passed through any of dozens of well known scrambling methods. The depiction of the number "28" alludes to the idea that there can actually be a library of scrambling methods, and the particular method used for this particular movie, or even for this particular frame, can change. The result is our classic master code signal or Snowy Image. In general, its brightness values are large and it would look very much like the snowy image on a television set tuned to a blank channel, but clearly it has been derived from an informative image 700, transformed through a scrambling 702. (Note: the splotchiness of the example picture is actually a rather poor depiction; it was a function of the crude tools available to the inventor).

This Master Snowy Image 704 is then the signal which is modulated by our N-bit identification word as outlined in other sections of the disclosure, the resulting modulated signal is then scaled down in brightness to the acceptable perceived noise level, and then added to the original frame to produce the distributable frame.

There are a variety of advantages and features that the method depicted in Fig. 13 affords. There are also variations of theme within this overall variation. Clearly, one advantage is that users can now use more intuitive and personalized methods for stamping and signing their work. Provided that the encryption/scrambling routines, 702, are indeed of a high security and not published or leaked, then even if a would-be pirate has knowledge of the logo image 700, they should not be able to use this knowledge to be able to sleuth the Master Snowy Image 704, and thus they should not be able to crack the system, as it were. On the other hand, simple encryption routines 702 may open the door for cracking the system. Another clear advantage of the method of Fig. 13 is the ability to place further information into the overall protective process. Strictly speaking, the information contained in the logo image 700 is not directly

- 42 -

carried in the final distributable frame. Said another way, and provided that the encryption/scrambling routine 702 has a straightforward and known decryption/descrambling method which tolerates bit truncation errors, it is generally impossible to fully re-create the image 700 based upon having the distributable frame, the N-bit identification code word, the brightness scaling factor used, and the number of the
5 decryption routine to be used. The reason that an exact recreation of the image 700 is impossible is due to the scaling operation itself and the concomitant bit truncation. For the present discussion, this whole issue is somewhat academic, however.

A variation on the theme of Fig. 13 is to actually place the N-bit identification code word directly into the logo image 700. In some sense this would be self-referential. Thus when we pull out our stored
10 logo image 700 it already contains visually what our identification word is, then we apply encryption routine #28 to this image, scale it down, then use this version to decode a suspect image using the techniques of this disclosure. The N bit word thus found should match the one contained in our logo image 700.

One desirable feature of the encryption/scrambling routine 702 might be (but is certainly not
15 required to be) that even given a small change in the input image 700, such as a single digit change of the frame number, there would be a huge visual change in the output scrambled master snowy image 704. Likewise, the actual scrambling routine may change as a function of frame numbers, or certain "seed" numbers typically used within pseudo-randomizing functions could change as a function of frame number. All manner of variations are thus possible, all helping to maintain high levels of security. Eventually,
20 engineering optimization considerations will begin to investigate the relationship between some of these randomizing methods, and how they all relate to maintaining acceptable signal strength levels through the process of transforming an uncompressed video stream into a compressed video stream such as with the MPEG compression methodologies.

Another desired feature of the encryption process 702 is that it should be informationally efficient,
25 i.e., that given any random input, it should be able to output an essentially spatially uniform noisy image with little to no residual spatial patterns beyond pure randomness. Any residual correlated patterns will contribute to inefficiency of encoding the N-bit identification word, as well as opening up further tools to would-be pirates to break the system.

Another feature of the method of Fig. 13 is that there is more intuitional appeal to using
30 recognizable symbols as part of a decoding system, which should then translate favorably in the essentially lay environment of a courtroom. It strengthens the simplicity of the coin flip vernacular mentioned elsewhere. Jury members or judges will better relate to an owner's logo as being a piece of the key of recognizing a suspect copy as being a knock-off.

It should also be mentioned that, strictly speaking, the logo image 700 does not need to be
35 randomized. The steps could equally apply straight to the logo image 700 directly. It is not entirely clear

to the inventor what practical goal this might have. A trivial extension of this concept to the case where $N=1$ is where, simply and easily, the logo image 700 is merely added to an original image at a very low brightness level. The inventor does not presume this trivial case to be at all a novelty. In many ways this is similar to the age old issue of subliminal advertising, where the low light level patterns added to an image are recognizable to the human eye/brain system and - supposedly - operating on the human brain at an unconscious level. By pointing out these trivial extensions of the current technology, hopefully there can arise further clarity which distinguishes applicant's novel principles in relation to such well known prior art techniques.

10

5-bit Abridged Alphanumeric Code Sets and Others

It is desirable in some applications for the N-bit identification word to actually signify names, companies, strange words, messages, and the like. Most of this disclosure focuses on using the N-bit identification word merely for high statistical security, indexed tracking codes, and other index based message carrying. The information carrying capacity of "invisible signatures" inside imagery and audio is somewhat limited, however, and thus it would be wise to use our N bits efficiently if we actually want to "spell out" alphanumeric items in the N-bit identification word.

One way to do this is to define, or to use an already existing, reduced bit (e.g. less than 8-bit ASCII) standardized codes for passing alphanumeric messages. This can help to satisfy this need on the part of some applications. For example, a simple alphanumeric code could be built on a 5-bit index table, where for example the letters V,X,Q, and Z are not included, but the digits 0 through 9 are included. In this way, a 100 bit identification word could carry with it 20 alphanumeric symbols. Another alternative is to use variable bit length codes such as the ones used in text compression routines (e.g. Huffman) whereby more frequently used symbols have shorter bit length codes and less frequently used symbols have longer bit lengths.

More on Detecting and Recognizing the N-bit Identification Word in Suspect Signals

Classically speaking, the detection of the N-bit identification word fits nicely into the old art of detecting known signals in noise. Noise in this last statement can be interpreted very broadly, even to the point where an image or audio track itself can be considered noise, relative to the need to detect the underlying signature signals. One of many references to this older art is the book Kassam, Saleem A., "Signal Detection in Non-Gaussian Noise," Springer-Verlag, 1988 (generally available at well stocked libraries, e.g. available at the U.S. Library of Congress by catalog number TK5102.5 .K357 1988,

To the best of this inventor's current understanding, none of the material

- 44 -

in this book is directly applicable to the issue of discovering the polarity of applicant's embedded signals, but the broader principles are indeed applicable.

In particular, section 1.2 "Basic Concepts of Hypothesis Testing" of Kassam's book lays out the basic concept of a binary hypothesis, assigning the value "1" to one hypothesis and the value "0" to the other hypothesis. The last paragraph of that section is also on point regarding the earlier-described embodiment, i.e., that the "0" hypothesis corresponds to "noise only" case, whereas the "1" corresponds to the presence of a signal in the observations. Applicant's use of true polarity is not like this, however, where now the "0" corresponds to the presence of an inverted signal rather than to "noise-only." Also in the present embodiment, the case of "noise-only" is effectively ignored, and that an identification process will either come up with our N-bit identification word or it will come up with "garbage."

The continued and inevitable engineering improvement in the detection of embedded code signals will undoubtedly borrow heavily from this generic field of known signal detection. A common and well-known technique in this field is the so-called "matched filter," which is incidentally discussed early in section 2 of the Kassam book. Many basic texts on signal processing include discussions on this method of signal detection. This is also known in some fields as correlation detection. Furthermore, when the phase or location of a known signal is known a priori, such as is often the case in applications of this technology, then the matched filter can often be reduced to a simple vector dot product between a suspect image and the embedded signal associated with an m'th bit plane in our N-bit identification word. This then represents yet another simple "detection algorithm" for taking a suspect image and producing a sequence of 1s and 0s with the intention of determining if that series corresponds to a pre-embedded N-bit identification word. In words, and with reference to Fig. 3, we run through the process steps up through and including the subtracting of the original image from the suspect, but the next step is merely to step through all N random independent signals and perform a simple vector dot product between these signals and the difference signal, and if that dot product is negative, assign a '0' and if that dot product is positive, assign a '1.' Careful analysis of this "one of many" algorithms will show its similarity to the traditional matched filter.

There are also some immediate improvements to the "matched filter" and "correlation-type" that can provide enhanced ability to properly detect very low level embedded code signals. Some of these improvements are derived from principles set forth in the Kassam book, others are generated by the inventor and the inventor has no knowledge of their being developed in other papers or works, but neither has the inventor done fully extensive searching for advanced signal detection techniques. One such technique is perhaps best exemplified by figure 3.5 in Kassam on page 79, wherein there are certain plots of the various locally optimum weighting coefficients which can apply to a general dot-product algorithmic approach to detection. In other words, rather than performing a simple dot product, each elemental multiplication operation in an overall dot product can be weighted based upon known *a priori* statistical

- 45 -

information about the difference signal itself, i.e., the signal within which the low level known signals are being sought. The interested reader who is not already familiar with these topics is encouraged to read chapter 3 of Kassam to gain a fuller understanding.

One principle which did not seem to be explicitly present in the Kassam book and which was developed rudimentarily by the inventor involves the exploitation of the magnitudes of the statistical properties of the known signal being sought relative to the magnitude of the statistical properties of the suspect signal as a whole. In particular, the problematic case seems to be where the embedded signals we are looking for are of much lower level than the noise and corruption present on a difference signal. Fig. 14 attempts to set the stage for the reasoning behind this approach. The top figure 720 contains a generic look at the differences in the histograms between a typical "problematic" difference signal, i.e., a difference signal which has a much higher overall energy than the embedded signals that may or may not be within it. The term "mean-removed" simply means that the means of both the difference signal and the embedded code signal have been removed, a common operation prior to performing a normalized dot product. The lower figure 722 then has a generally similar histogram plot of the derivatives of the two signals, or in the case of an image, the scalar gradients. From pure inspection it can be seen that a simple thresholding operation in the derivative transform domain, with a subsequent conversion back into the signal domain, will go a long way toward removing certain innate biases on the dot product "recognition algorithm" of a few paragraphs back. Thresholding here refers to the idea that if the absolute value of a difference signal derivative value exceeds some threshold, then it is replaced simply by that threshold value. The threshold value can be so chosen to contain most of the histogram of the embedded signal.

Another operation which can be of minor assistance in "alleviating" some of the bias effects in the dot product algorithm is the removal of the low order frequencies in the difference signal, i.e., running the difference signal through a high pass filter, where the cutoff frequency for the high pass filter is relatively near the origin (or DC) frequency.

Special Considerations for Recognizing Embedded Codes on Signals Which Have Been Compressed and Decompressed, or Alternatively, for Recognizing Embedded Codes Within Any Signal Which Has Undergone Some Known Process Which Creates Non-Uniform Error Sources

Long title for a basic concept. Some signal processing operations, such as compressing and decompressing an image, as with the JPEG/MPEG formats of image/video compression, create errors in some given transform domain which have certain correlations and structure. Using JPEG as an example, if a given image is compressed then decompressed at some high compression ratio, and that resulting image is then fourier transformed and compared to the fourier transform of the original uncompressed image, a definite pattern is clearly visible. This patterning is indicative of correlated error, i.e. error which can be to some extent quantified and predicted. The prediction of the grosser properties of this

- 46 -

correlated error can then be used to advantage in the heretofore-discussed methods of recognizing the embedded code signals within some suspect image which may have undergone either JPEG compression or any other operation which leaves these telltale correlated error signatures. The basic idea is that in areas where there are known higher levels of error, the value of the recognition methods is diminished relative to the areas with known lower levels of correlated errors. It is often possible to quantify the expected levels of error and use this quantification to appropriately weight the retransformed signal values. Using JPEG compression again as an example, a suspect signal can be fourier transformed, and the fourier space representation may clearly show the telltale box grid pattern. The fourier space signal can then be "spatially filtered" near the grid points, and this filtered representation can then be transformed back into its regular time or space domain to then be run through the recognition methods presented in this disclosure. Likewise, any signal processing method which creates non-uniform error sources can be transformed into the domain in which these error sources are non-uniform, the values at the high points of the error sources can be attenuated, and the thusly "filtered" signal can be transformed back into the time/space domain for standard recognition. Often this whole process will include the lengthy and arduous step of "characterizing" the typical correlated error behavior in order to "design" the appropriate filtering profiles.

"SIGNATURE CODES" and "INVISIBLE SIGNATURES"

Briefly and for the sake of clarity, the phrases and terms "signatures," "invisible signatures," and "signature codes" have been and will continue to be used to refer to the general techniques of this technology and often refer specifically to the composite embedded code signal as defined early on in this disclosure.

MORE DETAILS ON EMBEDDING SIGNATURE CODES INTO MOTION PICTURES

Just as there is a distinction made between the JPEG standards for compressing still images and the MPEG standards for compressed motion images, so too should there be distinctions made between placing invisible signatures into still images and placing signatures into motion images. As with the JPEG/MPEG distinction, it is not a matter of different foundations, it is the fact that with motion images a new dimension of engineering optimization opens up by the inclusion of time as a parameter. Any textbook dealing with MPEG will surely contain a section on how MPEG is (generally) not merely applying JPEG on a frame by frame basis. It will be the same with the application of the principles of this technology: generally speaking, the placement of invisible signatures into motion image sequences will not be simply independently placing invisible signatures into one frame after the next. A variety of time-based considerations come into play, some dealing with the psychophysics of motion image perception, others driven by simple cost engineering considerations.

- 47 -

One embodiment actually uses the MPEG compression standard as a piece of a solution. Other motion image compression schemes could equally well be used, be they already invented or yet to be invented. This example also utilizes the scrambled logo image approach to generating the master snowy image as depicted in Fig. 13 and discussed in the disclosure.

5 A "compressed master snowy image" is independently rendered as depicted in Fig. 15. "Rendered" refers to the generally well known technique in video, movie and animation production whereby an image or sequence of images is created by constructive techniques such as computer instructions or the drawing of animation cells by hand. Thus, "to render" a signature movie in this example is essentially to let either a computer create it as a digital file or to design some custom digital electronic circuitry to create it.

10 The overall goal of the procedure outlined in Fig. 15 is to apply the invisible signatures to the original movie 762 in such a way that the signatures do not degrade the commercial value of the movie, memorialized by the side-by-side viewing, 768, AND in such a way that the signature optimally survives through the MPEG compression and decompression process. As noted earlier, the use of the MPEG process in particular is an example of the generic process of compression. Also it should be noted that the example presented here has definite room for engineering variations. In particular, those practiced in the art of motion picture compression will appreciate the fact if we start out with two video streams A and B, and we compress A and B separately and combine their results, then the resultant video stream C will not generally be the same as if we pre-added the video streams A and B and compressed this resultant. Thus we have in general, e.g.:

$$\text{MPEG(A)} + \text{MPEG(B)} \neq \text{MPEG(A+B)}$$

25 where \neq is not equal to. This is somewhat an abstract notion to introduce at this point in the disclosure and will become more clear as Fig. 15 is discussed. The general idea, however, is that there will be a variety of algebras that can be used to optimize the pass-through of "invisible" signatures through compression procedures. Clearly, the same principles as depicted in Fig. 15 also work on still images and the JPEG or any other still image compression standard.

30 Turning now to the details of Fig. 15, we begin with the simple stepping through of all Z frames of a movie or video. For a two hour movie played at 30 frames per second, Z turns out to be $(30 \times 2 \times 60 \times 60)$ or 216,000. The inner loop of 700, 702 and 704 merely mimics Fig. 13's steps. The logo frame optionally can change during the stepping through frames. The two arrows emanating from the box 704 represent both the continuation of the loop 750 and the depositing of output frames into the rendered master Snowy Image 752.

- 48 -

To take a brief but potentially appropriate digression at this point, the use of the concept of a Markov process brings certain clarity to the discussion of optimizing the engineering implementation of the methods of Fig. 15. Briefly, a Markov process is one in which a sequence of events takes place and in general there is no memory between one step in the sequence and the next. In the context of Fig. 15 and a sequence of images, a Markovian sequence of images would be one in which there is no apparent or appreciable correlation between a given frame and the next. Imagine taking the set of all movies ever produced, stepping one frame at a time and selecting a random frame from a random movie to be inserted into an output movie, and then stepping through, say, one minute or 1800 of these frames. The resulting "movie" would be a fine example of a Markovian movie. One point of this discussion is that depending on how the logo frames are rendered and depending on how the encryption/scrambling step 702 is performed, the Master Snowy Movie 752 will exhibit some generally quantifiable degree of Markovian characteristics. The point of this point is that the compression procedure itself will be affected by this degree of Markovian nature and thus needs to be accounted for in designing the process of Fig. 15. Likewise, and only in general, even if a fully Markovian movie is created in the High Brightness Master Snowy Movie, 752, then the processing of compressing and decompressing that movie 752, represented as the MPEG box 754, will break down some of the Markovian nature of 752 and create at least a marginally non-Markovian compressed master Snowy Movie 756. This point will be utilized when the disclosure briefly discusses the idea of using multiple frames of a video stream in order to find a single N-bit identification word, that is, the same N-bit identification word may be embedded into several frames of a movie, and it is quite reasonable to use the information derived from those multiple frames to find that single N-bit identification word. The non-Markovian nature of 756 thus adds certain tools to reading and recognizing the invisible signatures. Enough of this tangent.

With the intent of pre-conditioning the ultimately utilized Master Snowy Movie 756, we now send the rendered High Brightness Master Snowy Movie 752 through both the MPEG compression AND decompression procedure 754. With the caveat previously discussed where it is acknowledged that the MPEG compression process is generally not distributive, the idea of the step 754 is to crudely segregate the initially rendered Snowy Movie 752 into two components, the component which survives the compression process 754 which is 756, and the component which does not survive, also crudely estimated using the difference operation 758 to produce the "Cheap Master Snowy Movie" 760. The reason use is made of the deliberately loose term "Cheap" is that we can later add this signature signal as well to a distributable movie, knowing that it probably won't survive common compression processes but that nevertheless it can provide "cheap" extra signature signal energy for applications or situations which will never experience compression. [Thus it is at least noted in Fig. 15]. Back to Fig. 15 proper, we now have a rough cut at signatures which we know have a higher likelihood of surviving intact through the compression process, and we use this "Compressed Master Snowy Movie" 756 to then go through this

procedure of being scaled down 764, added to the original movie 766, producing a candidate distributable movie 770, then compared to the original movie (768) to ensure that it meets whatever commercially viable criteria which have been set up (i.e. the acceptable perceived noise level). The arrow from the side-by-side step 768 back to the scale down step 764 corresponds quite directly to the "experiment visually..." step of Fig. 2, and the gain control 226 of Fig. 6. Those practiced in the art of image and audio information theory can recognize that the whole of Fig. 15 can be summarized as attempting to pre-condition the invisible signature signals in such a way that they are better able to withstand even quite appreciable compression. To reiterate a previously mentioned item as well, this idea equally applies to ANY such pre-identifiable process to which an image, and image sequence, or audio track might be subjected. This clearly includes the JPEG process on still images.

Additional Elements of the Realtime Encoder Circuitry

It should be noted that the method steps represented in Fig. 15, generally following from box 750 up through the creation of the compressed master snowy movie 756, could with certain modification be implemented in hardware. In particular, the overall analog noise source 206 in Fig. 6 could be replaced by such a hardware circuit. Likewise the steps and associated procedures depicted in Fig. 13 could be implemented in hardware and replace the analog noise source 206.

Recognition based on more than one frame: non-Markovian signatures

As noted in the digression on Markov and non-Markov sequences of images, it is pointed out once again that in such circumstances where the embedded invisible signature signals are non-Markovian in nature, i.e., that there is some correlation between the master snowy image of one frame to that of the next, AND furthermore that a single N-bit identification word is used across a range of frames and that the sequence of N-bit identification words associated with the sequence of frames is not Markovian in nature, then it is possible to utilize the data from several frames of a movie or video in order to recognize a single N-bit identification word. All of this is a fancy way of saying that the process of recognizing the invisible signatures should use as much information as is available, in this case translating to multiple frames of a motion image sequence.

30 HEADER VERIFICATION

The concept of the "header" on a digital image or audio file is a well established practice in the art. The top of Fig. 16 has a simplified look at the concept of the header, wherein a data file begins with generally a comprehensive set of information about the file as a whole, often including information about who the author or copyright holder of the data is, if there is a copyright holder at all. This header 800 is

- 50 -

then typically followed by the data itself 802, such as an audio stream, a digital image, a video stream, or compressed versions of any of these items. This is all exceedingly known and common in the industry.

One way in which the principles of this technology can be employed in the service of information integrity is generically depicted in the lower diagram of Fig. 16. In general, the N-bit identification word
5 can be used to essentially "wallpaper" a given simple message throughout an image (as depicted) or audio data stream, thereby reinforcing some message already contained in a traditional header. This is referred to as "header verification" in the title of this section. The thinking here is that less sophisticated would-be pirates and abusers can alter the information content of header information, and the more secure techniques of this technology can thus be used as checks on the veracity of header information. Provided that the
10 code message, such as "joe's image" in the header, matches the repeated message throughout an image, then a user obtaining the image can have some higher degree of confidence that no alteration of the header has taken place.

Likewise, the header can actually carry the N-bit identification word so that the fact that a given data set has been coded via the methods of this technology can be highlighted and the verification code
15 built right into the header. Naturally, this data file format has not been created yet since the principles of this technology are currently not being employed.

THE "BODIER": THE ABILITY TO LARGELY REPLACE A HEADER

Although all of the possible applications of the following aspect of applicant's technology are not
20 fully developed, it is nevertheless presented as a design alternative that may be important some day. The title of this section contains the silly phrase used to describe this possibility: the "bodier."

Whereas the previous section outlined how the N-bit identification word could "verify" information contained within the header of a digital file, there is also the prospect that these methods could completely replace the very concept of the header and place the information which is traditionally stored in
25 the header directly into the digital signal and empirical data itself.

This could be as simple as standardizing on, purely for example, a 96-bit (12 bytes) leader string on an otherwise entirely empirical data stream. This leader string would plain and simple contain the numeric length, in elemental data units, of the entire data file not including the leader string, and the number of bits of depth of a single data element (e.g. its number of grey levels or the number of discrete
30 signal levels of an audio signal). From there, universal codes as described in this specification would be used to read the N-bit identification word written directly within the empirical data. The length of the empirical data would need to be long enough to contain the full N bits. The N-bit word would effectively transmit what would otherwise be contained in a traditional header.

Fig. 17 depicts such a data format and calls it the "universal empirical data format." The leader
35 string 820 is comprised of the 64 bit string length 822 and the 32 bit data word size 824. The data stream

826 then immediately follows, and the information traditionally contained in the header but now contained directly in the data stream is represented as the attached dotted line 828. Another term used for this attached information is a "shadow channel" as also depicted in Fig. 17.

Yet another element that may need to be included in the leader string is some sort of complex
5 check sum bits which can verify that the whole of the data file is intact and unaltered. This is not included in Fig. 17.

MORE ON DISTRIBUTED UNIVERSAL CODE SYSTEMS: DYNAMIC CODES

One intriguing variation on the theme of universal codes is the possibility of the N-bit
10 identification word actually containing instructions which vary the operations of the universal code system itself. One of many examples is immediately in order: a data transmission is begun wherein a given block of audio data is fully transmitted, an N-bit identification word is read knowing that the first block of data used universal codes #145 out of a set of 500, say, and that part of the N-bit identification word thus found
15 is the instructions that the next block of data should be "analyzed" using the universal code set #411 rather than #145. In general, this technology can thus be used as a method for changing on the fly the actual decoding instructions themselves. Also in general, this ability to utilize "dynamic codes" should greatly increase the sophistication level of the data verification procedures and increase the economic viability of systems which are prone to less sophisticated thwarting by hackers and would-be pirates. The inventor
20 does not believe that the concept of dynamically changing decoding/decrypting instructions is novel *per se*, but the carrying of those instructions on the "shadow channel" of empirical data does appear to be novel to the best of the inventor's understanding. [Shadow channel was previously defined as yet another vernacular phrase encapsulating the more steganographic proper elements of this technology].

A variant on the theme of dynamic codes is the use of universal codes on systems which have *a priori* assigned knowledge of which codes to use when. One way to summarize this possibility is the idea
25 of "the daily password." The password in this example represents knowledge of which set of universal codes is currently operative, and these change depending on some set of application-specific circumstances. Presumably many applications would be continually updating the universal codes to ones which had never before been used, which is often the case with the traditional concept of the daily password. Part of a
30 currently transmitted N-bit identification word could be the passing on of the next day's password, for example. Though time might be the most common trigger events for the changing of passwords, there could be event based triggers as well.

SYMMETRIC PATTERNS AND NOISE PATTERNS: TOWARD A ROBUST UNIVERSAL CODING SYSTEM

- 52 -

The placement of identification patterns into images is certainly not new. Logos stamped into corners of images, subtle patterns such as true signatures or the wallpapering of the copyright circle-C symbol, and the watermark proper are all examples of placing patterns into images in order to signify ownership or to try to prevent illicit uses of the creative material.

5 What does appear to be novel is the approach of placing independent "carrier" patterns, which themselves are capable of being modulated with certain information, directly into images and audio for the purposes of transmission and discernment of said information, while effectively being imperceptible and/or unintelligible to a perceiving human. Steganographic solutions currently known to the inventor all place this information "directly" into empirical data (possibly first encrypted, then directly), whereas the methods
10 of this disclosure posit the creation of these (most-often) coextensive carrier signals, the modulation of those carrier signals with the information proper. THEN the direct application to the empirical data.

In extending these concepts one step further into the application arena of universal code systems, where a sending site transmits empirical data with a certain universal coding scheme employed and a receiving site analyzes said empirical data using the universal coding scheme, it would be advantageous to
15 take a closer look at the engineering considerations of such a system designed for the transmission of images or motion images, as opposed to audio. Said more clearly, the same type of analysis of a specific implementation such as is contained in Fig. 9 and its accompanying discussion on the universal codes in audio applications should as well be done on imagery (or two dimensional signals). This section is such an analysis and outline of a specific implementation of universal codes and it attempts to anticipate various
20 hurdles that such a method should clear.

The unifying theme of one implementation of a universal coding system for imagery and motion imagery is "symmetry." The idea driving this couldn't be more simple: a prophylactic against the use of image rotation as a means for less sophisticated pirates to bypass any given universal coding system. The guiding principle is that the universal coding system should easily be read no matter what rotational
25 orientation the subject imagery is in. These issues are quite common in the fields of optical character recognition and object recognition, and these fields should be consulted for further tools and tricks in furthering the engineering implementation of this technology. As usual, an immediate example is in order.

Digital Video And Internet Company XYZ has developed a delivery system of its product which relies on a non-symmetric universal coding which double checks incoming video to see if the individual
30 frames of video itself, the visual data, contain XYZ's own relatively high security internal signature codes using the methods of this technology. This works well and fine for many delivery situations, including their Internet tollgate which does not pass any material unless both the header information is verified AND the in-frame universal codes are found. However, another piece of their commercial network performs mundane routine monitoring on Internet channels to look for unauthorized transmission of their proprietary
35 creative property. They control the encryption procedures used, thus it is no problem for them to

unencrypt creative property, including headers, and perform straightforward checks. A pirate group that wants to traffic material on XYZ's network has determined how to modify the security features in XYZ's header information system, and they have furthermore discovered that by simply rotating imagery by 10 or 20 degrees, and transmitting it over XYZ's network, the network doesn't recognize the codes and therefore does not flag illicit uses of their material, and the receiver of the pirate's rotated material simply unrotates it.

Summarizing this last example via logical categories, the non-symmetric universal codes are quite acceptable for the "enablement of authorized action based on the finding of the codes," whereas it can be somewhat easily by-passed in the case of "random monitoring (policing) for the presence of codes." [Bear in mind that the non-symmetric universal codes may very well catch 90% of illicit uses, i.e. 90% of the illicit users wouldn't bother even going to the simple by-pass of rotation.] To address this latter category, the use of quasi-rotationally symmetric universal codes is called for. "Quasi" derives from the age old squaring the circle issue, in this instance translating into not quite being able to represent a full incrementally rotational symmetric 2-D object on a square grid of pixels. Furthermore, basic considerations must be made for scale/magnification changes of the universal codes. It is understood that the monitoring process must be performed when the monitored visual material is in the "perceptual" domain, i.e. when it has been unencrypted or unscrambled and in the form with which it is (or would be) presented to a human viewer. Would-be pirates could attempt to use other simple visual scrambling and unscrambling techniques, and tools could be developed to monitor for these telltale scrambled signals. Said another way, would-be pirates would then look to transform visual material out of the perceptual domain, pass by a monitoring point, and then transform the material back into the perceptual domain; tools other than the monitoring for universal codes would need to be used in such scenarios. The monitoring discussed here therefore applies to applications where monitoring can be performed in the perceptual domain, such as when it is actually sent to viewing equipment.

The "ring" is the only full rotationally symmetric two dimensional object. The "disk" can be seen as a simple finite series of concentric and perfectly abutted rings having width along their radial axis. Thus, the "ring" needs to be the starting point from which a more robust universal code standard for images is found. The ring also will fit nicely into the issue of scale/magnification changes, where the radius of a ring is a single parameter to keep track of and account for. Another property of the ring is that even the case where differential scale changes are made to different spatial axes in an image, and the ring turns into an oval, many of the smooth and quasi-symmetric properties that any automated monitoring system will be looking for are generally maintained. Likewise, appreciable geometric distortion of any image will clearly distort rings but they can still maintain gross symmetric properties. Hopefully, more pedestrian methods such as simply "viewing" imagery will be able to detect attempted illicit piracy in these regards, especially when such lengths are taken to by-pass the universal coding system.

Rings to Knots

Having discovered the ring as the only ideal symmetric pattern upon whose foundation a full rotationally robust universal coding system can be built, we must turn this basic pattern into something functional, something which can carry information, can be read by computers and other instrumentation, can survive simple transformations and corruptions, and can give rise to reasonably high levels of security (probably not unbreakable, as the section on universal codes explained) in order to keep the economics of subversion as a simple incremental cost item.

One embodiment of the "ring-based" universal codes is what the inventor refers to as "knot patterns" or simply "knots," in deference to woven Celtic knot patterns which were later refined and exalted in the works of Leonardo Da Vinci (e.g. Mona Lisa, or his knot engravings). Some rumors have it that these drawings of knots were indeed steganographic in nature, i.e. conveying messages and signatures; all the more appropriate. Figs. 18 and 19 explore some of the fundamental properties of these knots.

Two simple examples of knot patterns are depicted by the supra-radial knots, 850 and the radial knots 852. The names of these types are based on the central symmetry point of the splayed rings and whether the constituent rings intersect this point, are fully outside it, or in the case of sub-radial knots the central point would be inside a constituent circle. The examples of 850 and 852 clearly show a symmetrical arrangement of 8 rings or circles. "Rings" is the more appropriate term, as discussed above, in that this term explicitly acknowledges the width of the rings along the radial axis of the ring. It is each of the individual rings in the knot patterns 850 and 852 which will be the carrier signal for a single associated bit plane in our N-bit identification word. Thus, the knot patterns 850 and 852 each are an 8-bit carrier of information. Specifically, assuming now that the knot patterns 850 and 852 are luminous rings on a black background, then the "addition" of a luminous ring to an independent source image could represent a "1" and the "subtraction" of a luminous ring from an independent source image could represent a "0." The application of this simple encoding scheme could then be replicated over and over as in Fig. 19 and its mosaic of knot patterns, with the ultimate step of adding a scaled down version of this encoded (modulated) knot mosaic directly and coextensively to the original image, with the resultant being the distributable image which has been encoded via this universal symmetric coding method. It remains to communicate to a decoding system which ring is the least significant bit in our N-bit identification word and which is the most significant. One such method is to make a slightly ascending scale of radii values (of the individual rings) from the LSB to the MSB. Another is to merely make the MSB, say, 10% larger radius than all the others and to pre-assign counterclockwise as the order with which the remaining bits fall out. Yet another is to put some simple hash mark inside one and only one circle. In other words, there are a variety of ways with which the bit order of the rings can be encoded in these knot patterns.

- 55 -

A procedure for, first, checking for the mere existence of these knot patterns and, second, for reading of the N-bit identification word, is as follows. A suspect image is first fourier transformed via the extremely common 2D FFT computer procedure. Assuming that we don't know the exact scale of the knot patterns, i.e., we don't know the radius of an elemental ring of the knot pattern in the units of pixels, and that we don't know the exact rotational state of a knot pattern, we merely inspect (via basic automated pattern recognition methods) the resulting magnitude of the Fourier transform of the original image for telltale ripple patterns (concentric low amplitude sinusoidal rings on top of the spatial frequency profile of a source image). The periodicity of these rings, along with the spacing of the rings, will inform us that the universal knot patterns are or are not likely present, and their scale in pixels. Classical small signal detection methods can be applied to this problem just as they can to the other detection methodologies of this disclosure. Common spatial filtering can then be applied to the fourier transformed suspect image, where the spatial filter to be used would pass all spatial frequencies which are on the crests of the concentric circles and block all other spatial frequencies. The resulting filtered image would be fourier transformed out of the spatial frequency domain back into the image space domain, and almost by visual inspection the inversion or non-inversion of the luminous rings could be detected, along with identification of the MSB or LSB ring, and the (in this case 8) N-bit identification code word could be found. Clearly, a pattern recognition procedure could perform this decoding step as well.

The preceding discussion and the method it describes has certain practical disadvantages and shortcomings which will now be discussed and improved upon. The basic method was presented in a simple-minded fashion in order to communicate the basic principles involved.

Let's enumerate a few of the practical difficulties of the above described universal coding system using the knot patterns. For one (1), the ring patterns are somewhat inefficient in their "covering" of the full image space and in using all of the information carrying capacity of an image extent. Second (2), the ring patterns themselves will almost need to be more visible to the eye if they are applied, say, in a straightforward additive way to an 8-bit black and white image. Next (3), the "8" rings of Fig. 18, 850 and 852, is a rather low number, and moreover, there is a 22 and one half degree rotation which could be applied to the figures which the recognition methods would need to contend with (360 divided by 8 divided by 2). Next (4), strict overlapping of rings would produce highly condensed areas where the added and subtracted brightness could become quite appreciable. Next (5), the 2D FFT routine used in the decoding is notoriously computationally cumbersome as well as some of the pattern recognition methods alluded to. Finally (6), though this heretofore described form of universal coding does not pretend to have ultra-high security in the classical sense of top security communications systems, it would nevertheless be advantageous to add certain security features which would be inexpensive to implement in hardware and software systems which at the same time would increase the cost of would-be pirates attempting to thwart the system, and increase the necessary sophistication level of those pirates, to the point that a would-be

pirate would have to go so far out of their way to thwart the system that willfulness would be easily proven and hopefully subject to stiff criminal liability and penalty (such as the creation and distribution of tools which strip creative property of these knot pattern codes).

All of these items can be addressed and should continue to be refined upon in any engineering
5 — implementation of the principles of the technology. This disclosure addresses these items with reference to the following embodiments.

Beginning with item number 3, that there are only 8 rings represented in Fig. 18 is simply remedied by increasing the number of rings. The number of rings that any given application will utilize is clearly a function of the application. The trade-offs include but are not limited to: on the side which
10 argues to limit the number of rings utilized, there will ultimately be more signal energy per ring (per visibility) if there are less rings; the rings will be less crowded so that there discernment via automated recognition methods will be facilitated; and in general since they are less crowded, the full knot pattern can be contained using a smaller overall pixel extent, e.g. a 30 pixel diameter region of image rather than a 100 pixel diameter region. The arguments to increase the number of rings include: the desire to transmit
15 more information, such as ascii information, serial numbers, access codes, allowed use codes and index numbers, history information, etc.; another key advantage of having more rings is that the rotation of the knot pattern back into itself is reduced, thereby allowing the recognition methods to deal with a smaller range of rotation angles (e.g., 64 rings will have a maximum rotational displacement of just under 3 degrees, i.e. maximally dissimilar to its original pattern, where a rotation of about 5 and one half degrees
20 brings the knot pattern back into its initial alignment; the need to distinguish the MSB/LSB and the bit plane order is better seen in this example as well). It is anticipated that most practical applications will choose between 16 and 128 rings, corresponding to $N=16$ to $N=128$ for the choice of the number of bits in the N-bit identification code word. The range of this choice would somewhat correlate to the overall radius, in pixels, allotted to an elemental knot pattern such as 850 or 852.

25 Addressing the practical difficulty item number 4, that of the condensation of rings patterns at some points in the image and lack of ring patterns in others (which is very similar, but still distinct from, item 1, the inefficient covering), the following improvement can be applied. Fig. 18 shows an example of a key feature of a "knot" (as opposed to a pattern of rings) in that where patterns would supposedly intersect, a virtual third dimension is posited whereby one strand of the knot takes precedence over another
30 strand in some predefined way; see item 854. In the terms of imagery, the brightness or dimness of a given intersection point in the knot patterns would be "assigned" to one and only one strand, even in areas where more than two strands overlap. The idea here is then extended, 864, to how rules about this assignment should be carried out in some rotationally symmetric manner. For example, a rule would be that, travelling clockwise, an incoming strand to a loop would be "behind" an outgoing strand. Clearly
35 there are a multitude of variations which could be applied to these rules, many which would critically

depend on the geometry of the knot patterns chosen. Other issues involved will probably be that the finite width, and moreover the brightness profile of the width along the normal axis to the direction of a strand, will all play a role in the rules of brightness assignment to any given pixel underlying the knot patterns.

A major improvement to the nominal knot pattern system previously described directly addresses practical difficulties (1), the inefficient covering, (2) the unwanted visibility of the rings, and (6) the need for higher levels of security. This improvement also indirectly address item (4) the overlapping issue, which has been discussed in the last paragraph. This major improvement is the following: just prior to the step where the mosaic of the encoded knot patterns is added to an original image to produce a distributable image, the mosaic of encoded knot patterns, 866, is spatially filtered (using common 2D FFT techniques) by a standardized and (generally smoothly) random phase-only spatial filter. It is very important to note that this phase-only filter is itself fully rotationally symmetric within the spatial frequency domain, i.e. its filtering effects are fully rotationally symmetric. The effect of this phase-only filter on an individual luminous ring is to transform it into a smoothly varying pattern of concentric rings, not totally dissimilar to the pattern on water several instances after a pebble is dropped in, only that the wave patterns are somewhat random in the case of this phase-only filter rather than the uniform periodicity of a pebble wave pattern. Fig. 20 attempts to give a rough (i.e. non-greyscale) depiction of these phase-only filtered ring patterns. The top figure of Fig. 20 is a cross section of a typical brightness contour/profile of one of these phase-only filtered ring patterns. Referenced in the figure is the nominal location of the pre-filtered outer ring center, 870. The center of an individual ring, 872, is referenced as the point around which the brightness profile is rotated in order to fully describe the two dimensional brightness distribution of one of these filtered patterns. Yet another rough attempt to communicate the characteristics of the filtered ring is depicted as 876, a crude greyscale image of the filtered ring. This phase-only filtered ring, 876 will can be referred to as a random ripple pattern.

Not depicted in Fig. 20 is the composite effects of phase-only filtering on the knot patterns of Fig. 18, or on the mosaic of knot patterns 866 in Fig. 19. Each of the individual rings in the knot patterns 850 or 852 will give rise to a 2D brightness pattern of the type 876, and together they form a rather complicated brightness pattern. Realizing that the encoding of the rings is done by making it luminous (1) or "anti-luminous" (0), the resulting phase-only filtered knot patterns begin to take on subtle characteristics which no longer make direct sense to the human eye, but which are still readily discernable to a computer especially after the phase-only filtering is inverse filtered reproducing the original rings patterns.

Returning now to Fig. 19, we can imagine that an 8-bit identification word has been encoded on the knot patterns and the knot patterns phase-only filtered. The resulting brightness distribution would be a rich tapestry of overlapping wave patterns which would have a certain beauty, but would not be readily intelligible to the eye/brain. [An exception to this might draw from the lore of the South Pacific Island communities, where it is said that sea travellers have learned the subtle art of reading small and multiply

complex ocean wave patterns, generated by diffracted and reflected ocean waves off of intervening islands, as a primary navigational tool.] For want of a better term, the resulting mosaic of filtered knot patterns (derived from 866) can be called the encoded knot tapestry or just the knot tapestry. Some basic properties of this knot tapestry are that it retains the basic rotational symmetry of its generator mosaic, it is generally unintelligible to the eye/brain, thus raising it a notch on the sophistication level of reverse engineering, it is more efficient at using the available information content of a grid of pixels (more on this in the next section), and if the basic knot concepts 854 and 864 are utilized, it will not give rise to local "hot spots" where the signal level becomes unduly condensed and hence objectionably visible to a viewer.

The basic decoding process previously described would now need the additional step of inverse filtering the phase-only filter used in the encoding process. This inverse filtering is quite well known in the image processing industry. Provided that the scale of the knot patterns are known *a priori*, the inverse filtering is straightforward. If on the other hand the scale of the knot patterns is not known, then an additional step of discovering this scale is in order. One such method of discovering the scale of the knot patterns is to iteratively apply the inverse phase-only filter to variously scaled version of an image being decoded, searching for which scale-version begins to exhibit noticeable knot patterning. A common search algorithm such as the simplex method could be used in order to accurately discover the scale of the patterns. The field of object recognition should also be consulted, under the general topic of unknown-scale object detection.

An additional point about the efficiency with which the knot tapestry covers the image pixel grid is in order. Most applications of the knot tapestry method of universal image coding will posit the application of the fully encoded tapestry (i.e. the tapestry which has the N-bit identification word embedded) at a relative low brightness level into the source image. In real terms, the brightness scale of the encoded tapestry will vary from, for example, -5 grey scale values to 5 grey scale values in a typical 256 grey scale image, where the preponderance of values will be within -2 and 2. This brings up the purely practical matter that the knot tapestry will be subject to appreciable bit truncation error. Put as an example, imagine a constructed knot tapestry nicely utilizing a full 256 grey level image, then scaling this down by a factor of 20 in brightness including the bit truncation step, then rescaling this truncated version back up in brightness by the same factor of 20, then inverse phase-only filtering the resultant. The resulting knot pattern mosaic will be a noticeably degraded version of the original knot pattern mosaic. The point of bringing all of this up is the following: it will be a simply defined, but indeed challenging, engineering task to select the various free parameters of design in the implementation of the knot tapestry method, the end goal being to pass a maximum amount of information about the N-bit identification word within some pre-defined visibility tolerance of the knot tapestry. The free parameters include but would not be fully limited to: the radius of the elemental ring in pixels, N or the number of rings, the distance in pixels from the center of a knot pattern to the center of an elemental ring, the packing criteria and

- 59 -

distances of one knot pattern with the others, the rules for strand weaving, and the forms and types of phase-only filters to be used on the knot mosaics. It would be desirable to feed such parameters into a computer optimization routine which could assist in their selection. Even this would begin surely as more of an art than a science due to the many non-linear free parameters involved.

5 A side note on the use of phase-only filtering is that it can assist in the detection of the ring patterns. It does so in that the inverse filtering of the decoding process tends to "blur" the underlying source image upon which the knot tapestry is added, while at the same time "bringing into focus" the ring patterns. Without the blurring of the source image, the emerging ring patterns would have a harder time "competing" with the sharp features of typical images. The decoding procedure should also utilize the
10 gradient thresholding method described in another section. Briefly, this is the method where if it is known that a source signal is much larger in brightness than our signature signals, then an image being decoded can have higher gradient areas thresholded in the service of increasing the signal level of the signature signals relative to the source signal.

As for the other practical difficulty mentioned earlier, item (5) which deals with the relative
15 computational overhead of the 2D FFT routine and of typical pattern recognition routines, the first remedy here posited but not filled is to find a simpler way of quickly recognizing and decoding the polarity of the ring brightnesses than that of using the 2D FFT. Barring this, it can be seen that if the pixel extent of an individual knot pattern (850 or 852) is, for example, 50 pixels in diameter, then a simple 64 by 64 pixel 2D FFT on some section of an image may be more than sufficient to discern the N-bit identification word
20 as previously described. The idea would be to use the smallest image region necessary, as opposed to being required to utilize an entire image, to discern the N-bit identification word.

Another note is that those practitioners in the science of image processing will recognize that instead of beginning the discussion on the knot tapestry with the utilization of rings, we could have instead jumped right to the use of 2D brightness distribution patterns 876, QUA bases functions. The use of the
25 "ring" terminology as the baseline technology is partly didactic, as is appropriate for patent disclosures in any event. What is more important, perhaps, is that the use of true "rings" in the decoding process, post-inverse filtering, is probably the simplest form to input into typical pattern recognition routines.

Neural Network Decoders

30 Those skilled in the signal processing art will recognize that computers employing neural network architectures are well suited to the pattern recognition and detection-of-small-signal-in-noise issues posed by the present technology. While a complete discourse on these topics is beyond the scope of this specification, the interested reader is referred to, e.g., Cherkassky, V., "From Statistics to Neural Networks: Theory & Pattern Recognition Applications," Springer-Verlag, 1994; Masters, T., "Signal & Image Processing with Neural Networks: C Sourcebook," Wiley, 1994; Guyon, I., "Advances in Pattern
35

Recognition Systems Using Neural Networks, " World Scientific Publishers, 1994; Nigrin, A., "Neural Networks for Pattern Recognition," MIT Press, 1993; Cichoki, A., "Neural Networks for Optimization & Signal Processing," Wiley, 1993; and Chen, C., "Neural Networks for Pattern Recognition & Their Applications," World Scientific Publishers, 1991.

5

2D UNIVERSAL CODES II : SIMPLE SCAN LINE IMPLEMENTATION OF THE ONE DIMENSIONAL CASE

The above section on rings, knots and tapestries certainly has its beauty, but some of the steps involved may have enough complexity that practical implementations may be too costly for certain applications. A poor cousin the concept of rings and well-designed symmetry is to simply utilize the basic concepts presented in connection with Fig. 9 and the audio signal, and apply them to two dimensional signals such as images, but to do so in a manner where, for example, each scan line in an image has a random starting point on, for example, a 1000 pixel long universal noise signal. It would then be incumbent upon recognition software and hardware to interrogate imagery across the full range of rotational states and scale factors to "find" the existence of these universal codes.

15

THE UNIVERSAL COMMERCIAL COPYRIGHT (UCC) IMAGE, AUDIO, AND VIDEO FILE FORMATS

It is as well known as it is regretted that there exist a plethora of file format standards (and not-so-standards) for digital images, digital audio, and digital video. These standards have generally been formed within specific industries and applications, and as the usage and exchange of creative digital material proliferated, the various file formats slugged it out in cross-disciplinary arenas, where today we find a *de facto* histogram of devotees and users of the various favorite formats. The JPEG, MPEG standards for formatting and compression are only slight exceptions it would seem, where some concerted cross-industry collaboration came into play.

25

The cry for a simple universal standard file format for audio/visual data is as old as the hills. The cry for the protection of such material is older still. With all due respect to the innate difficulties attendant upon the creation of a universal format, and with all due respect to the pretentiousness of outlining such a plan within a patent disclosure, the inventor does believe that these methods can serve perhaps as well as anything for being the foundation upon which an accepted world-wide "universal commercial copyright" format is built. Practitioners know that such animals are not built by proclamation, but through the efficient meeting of broad needs, tenacity, and luck. More germane to the purposes of this disclosure is the fact that the application of this technology would benefit if it could become a central piece within an industry standard file format. The use of universal codes in particular could be specified within such a

30

- 61 -

standard. The fullest expression of the commercial usage of this technology comes from the knowledge that the invisible signing is taking place and the confidence that instills in copyright holders.

The following is a list of reasons that the principles of this technology could serve as the catalyst for such a standard: (1) Few if any technical developments have so isolated and so pointedly addressed the issue of broad-brush protection of empirical data and audio/visual material; (2) All previous file formats have treated the information about the data, and the data itself, as two separate and physically distinct entities, whereas the methods of this technology can combine the two into one physical entity; (3) The mass scale application of the principles of this technology will require substantial standardization work in the first place, including integration with the years-to-come improvements in compression technologies, so the standards infrastructure will exist by default; (4) the growth of multimedia has created a generic class of data called "content," which includes text, images, sound, and graphics, arguing for higher and higher levels of "content standards"; and (5) marrying copyright protection technology and security features directly into a file format standard is long overdue.

Elements of a universal standard would certainly include the mirroring aspects of the header verification methods, where header information is verified by signature codes directly within data. Also, a universal standard would outline how hybrid uses of fully private codes and public codes would commingle. Thus, if the public codes were "stripped" by sophisticated pirates, the private codes would remain intact. A universal standard would specify how invisible signatures would evolve as digital images and audio evolve. Thus, when a given image is created based on several source images, the standard would specify how and when the old signatures would be removed and replaced by new signatures, and if the header would keep track of these evolutions and if the signatures themselves would keep some kind of record.

PIXELS VS. BUMPS

Most of the disclosure focuses on pixels being the basic carriers of the N-bit identification word. The section discussing the use of a single "master code signal" went so far as to essentially "assign" each and every pixel to a unique bit plane in the N-bit identification word.

For many applications, with one exemplar being that of ink based printing at 300 dots per inch resolution, what was once a pixel in a pristine digital image file becomes effectively a blob (e.g. of dithered ink on a piece of paper). Often the isolated information carrying capacity of the original pixel becomes compromised by neighboring pixels spilling over into the geometrically defined space of the original pixel. Those practiced in the art will recognize this as simple spatial filtering and various forms of blurring.

In such circumstances it may be more advantageous to assign a certain highly local group of pixels to a unique bit plane in the N-bit identification word, rather than merely a single pixel. The end goal is

simply to pre-concentrate more of the signature signal energy into the lower frequencies, realizing that most practical implementations quickly strip or mitigate higher frequencies.

A simple-minded approach would be to assign a 2 by 2 block of pixels all to be modulated with the same ultimate signature grey value, rather than modulating a single assigned pixel. A more fancy approach is depicted in Figs. 21A and 21B, where an array of pixel groups is depicted. This is a specific example of a large class of configurations. The idea is that now a certain small region of pixels is associated with a given unique bit plane in the N-bit identification word, and that this grouping actually shares pixels between bit planes (though it doesn't necessary have to share pixels, as in the case of a 2x2 block of pixels above).

Depicted in Figs. 21A and 21B is a 3x3 array of pixels with an example normalized weighting (normalized --> the weights add up to 1). The methods of this technology now operate on this elementary "bump," as a unit, rather than on a single pixel. It can be seen that in this example there is a fourfold decrease in the number of master code values that need to be stored, due to the spreading out of the signature signal. Applications of this "bump approach" to placing in invisible signatures include any application which will experience *a priori* known high amounts of blurring, where proper identification is still desired even after this heavy blurring.

MORE ON THE STEGANOGRAPHIC USES OF THIS TECHNOLOGY

As mentioned in the initial sections of the disclosure, steganography as an art and as a science is a generic prior art to this technology. Putting the shoe on the other foot now, and as already doubtless apparent to the reader who has ventured thus far, the methods of this technology can be used as a novel method for performing steganography. (Indeed, all of the discussion thus far may be regarded as exploring various forms and implementations of steganography.)

In the present section, we shall consider steganography as the need to pass a message from point A to point B, where that message is essentially hidden within generally independent empirical data. As anyone in the industry of telecommunications can attest to, the range of purposes for passing messages is quite broad. Presumably there would be some extra need, beyond pure hobby, to place messages into empirical data and empirical signals, rather than sending those messages via any number of conventional and straightforward channels. Past literature and product propaganda within steganography posits that such an extra need, among others, might be the desire to hide the fact that a message is even being sent. Another possible need is that a conventional communications channel is not available directly or is cost prohibitive, assuming, that is, that a sender of messages can "transmit" their encoded empirical data somehow. This disclosure includes by reference all previous discussions on the myriad uses to which steganography might apply, while adding the following uses which the inventor has not previously seen described.

- 63 -

The first such use is very simple. It is the need to carry messages about the empirical data within which the message is carried. The little joke is that now the media is truly the message, though it would be next to impossible that some previous steganographer hasn't already exploited this joke. Some of the discussion on placing information about the empirical data directly inside that empirical data was already covered in the section on replacing the header and the concept of the "bodier." This section expands upon that section somewhat.

The advantages of placing a message about empirical data directly in that data is that now only one class of data object is present rather than the previous two classes. In any two class system, there is the risk of the two classes becoming disassociated, or one class corrupted without the other knowing about it. A concrete example here is what the inventor refers to as "device independent instructions."

There exist zillions of machine data formats and data file formats. This plethora of formats has been notorious in its power to impede progress toward universal data exchange and having one machine do the same thing that another machine can do. The instructions that an originator might put into a second class of data (say the header) may not at all be compatible with a machine which is intended to recognize these instructions. If format conversions have taken place, it is also possible that critical instructions have been stripped along the way, or garbled. The improvements disclosed here can be used as a way to "seal in" certain instructions directly into empirical data in such a way that all that is needed by a reading machine to recognize instructions and messages is to perform a standardized "recognition algorithm" on the empirical data (providing of course that the machine can at the very least "read" the empirical data properly). All machines could implement this algorithm any old way they choose, using any compilers or internal data formats that they want.

Implementation of this device independent instruction method would generally not be concerned over the issue of piracy or illicit removal of the sealed in messages. Presumably, the embedded messages and instructions would be a central valuable component in the basic value and functioning of the material.

Another example of a kind of steganographic use of the technology is the embedding of universal use codes for the benefit of a user community. The "message" being passed could be simply a registered serial number identifying ownership to users who wish to legitimately use and pay for the empirical information. The serial number could index into a vast registry of creative property, containing the name or names of the owners, pricing information, billing information, and the like. The "message" could also be the clearance of free and public use for some given material. Similar ownership identification and use indexing can be achieved in two class data structure methods such as a header, but the use of the single class system of this technology may offer certain advantages over the two class system in that the single class system does not care about file format conversion, header compatibilities, internal data format issues, header/body archiving issues, and media transformations.

35

Fully Exact Steganography

Prior art steganographic methods currently known to the inventor generally involve fully deterministic or "exact" prescriptions for passing a message. Another way to say this is that it is a basic assumption that for a given message to be passed correctly in its entirety, the receiver of the information
5 needs to receive the exact digital data file sent by the sender, tolerating no bit errors or "loss" of data. By definition, "lossy" compression and decompression on empirical signals defeat such steganographic methods. (Prior art, such as the previously noted Komatsu work, are the exceptions here.)

The principles of this technology can also be utilized as an exact form of steganography proper. It is suggested that such exact forms of steganography, whether those of prior art or those of this
10 technology, be combined with the relatively recent art of the "digital signature" and/or the DSS (digital signature standard) in such a way that a receiver of a given empirical data file can first verify that not one single bit of information has been altered in the received file, and thus verify that the contained exact steganographic message has not been altered.

The simplest way to use the principles of this technology in an exact steganographic system is to
15 utilize the previously discussed "designed" master noise scheme wherein the master snowy code is not allowed to contain zeros. Both a sender and a receiver of information would need access to BOTH the master snowy code signal AND the original unencoded original signal. The receiver of the encoded signal merely subtracts the original signal giving the difference signal and the techniques of simple polarity checking between the difference signal and the master snowy code signal, data sample to data sample,
20 producing a the passed message a single bit at a time. Presumably data samples with values near the "rails" of the grey value range would be skipped (such as the values 0,1,254 and 255 in 8-bit depth empirical data).

Statistical Steganography

25 The need for the receiver of a steganographic embedded data file to have access to the original signal can be removed by turning to what the inventor refers to as "statistical steganography." In this approach, the methods of this technology are applied as simple *a priori* rules governing the reading of an empirical data set searching for an embedded message. This method also could make good use of it combination with prior art methods of verifying the integrity of a data file, such as with the DSS. (See,
30 e.g., Walton, "Image Authentication for a Slippery New Age," Dr. Dobb's Journal, April, 1995, p. 18 for methods of verifying the sample-by-sample, bit-by-bit, integrity of a digital image.)

Statistical steganography posits that a sender and receiver both have access to the same master snowy code signal. This signal can be entirely random and securely transmitted to both parties, or generated by a shared and securely transmitted lower order key which generates a larger quasi-random
35 master snowy code signal. It is *a priori* defined that 16 bit chunks of a message will be passed within

contiguous 1024 sample blocks of empirical data, and that the receiver will use dot product decoding methods as outlined in this disclosure. The sender of the information pre-checks that the dot product approach indeed produces the accurate 16 bit values (that is, the sender pre-checks that the cross-talk between the carrier image and the message signal is not such that the dot product operation will produce an unwanted inversion of any of the 16 bits). Some fixed number of 1024 sample blocks are transmitted and the same number times 16 bits of message is therefore transmitted. DSS techniques can be used to verify the integrity of a message when the transmitted data is known to only exist in digital form, whereas internal checksum and error correcting codes can be transmitted in situations where the data may be subject to change and transformation in its transmission. In this latter case, it is best to have longer blocks of samples for any given message content size (such as 10K samples for a 16 bit message chunk, purely as an example).

Continuing for a moment on the topic of error correcting steganography, it will be recognized that many decoding techniques disclosed herein operate on the principle of distinguishing pixels (or bumps) which have been augmented by the encoded data from those that have been diminished by the encoded data. The distinguishing of these positive and negative cases becomes increasingly difficult as the delta values (e.g. the difference between an encoded pixel and the corresponding original pixel) approach zero.

An analogous situation arises in certain modem transmission techniques, wherein an ambiguous middle ground separates the two desired signal states (e.g. ± 1). Errors deriving from incorrect interpretation of this middle ground are sometimes termed "soft errors." Principles from modem technology, and other technologies where such problems arise, can likewise be applied to mitigation of such errors in the present context.

One approach is to weight the "confidence" of each delta determination. If the pixel (bump) clearly reflects one state or the other (e.g. ± 1), its "confidence" is said to be high, and it is given a proportionately greater weighting. Conversely, if the pixel (bump) is relatively ambiguous in its interpretation, its confidence is commensurately lower, and it is given a proportionately lesser weighting. By weighting the data from each pixel (bump) in accordance with its confidence value, the effects of soft errors can be greatly reduced.

The foregoing procedure, while theoretically simple, relies on weightings which are best determined empirically. Accordingly, such an approach is not necessarily straightforward.

An alternative approach is to assign confidence values not to interpretations of individual pixels, but rather to determination of bit values -- either from an image excerpt, or across the entire image. In such an arrangement, each decoded message bit is given a confidence value depending on the ambiguity (or not) of the image statistics by which its value was determined.

- 66 -

Such confidence weighting can also be used as a helpful adjunct with other error detecting/correcting schemes. For example, in known error correcting polynomials, the above-detailed weighting parameters can be used to further hone polynomial-based discernment of an error's location.

5 THE "NOISE" IN VECTOR GRAPHICS AND VERY-LOW-ORDER INDEXED GRAPHICS

The methods of this disclosure generally posit the existence of "empirical signals," which is another way of saying signals which have noise contained within them almost by definition. There are two classes of 2 dimensional graphics which are not generally considered to have noise inherent in them: vector graphics and certain indexed bit-mapped graphics. Vector graphics and vector graphic files are generally files which contain exact instructions for how a computer or printer draws lines, curves and shapes. A change of even one bit value in such a file might change a circle to a square, as a very crude example. In other words, there is generally no "inherent noise" to exploit within these files. Indexed bit-mapped graphics refers to images which are composed of generally a small number of colors or grey values, such as 16 in the early CGA displays on PC computers. Such "very-low-order" bit-mapped images usually display graphics and cartoons, rather than being used in the attempted display of a digital image taken with a camera of the natural world. These types of very-low-order bit-mapped graphics also are generally not considered to contain "noise" in the classic sense of that term. The exception is where indexed graphic files do indeed attempt to depict natural imagery, such as with the GIF (graphic interchange format of Compuserve), where the concept of "noise" is still quite valid and the principles of this technology still quite valid. These latter forms often use dithering (similar to pointillist paintings and color newspaper print) to achieve near lifelike imagery.

This section concerns this class of 2 dimensional graphics which traditionally do not contain "noise." This section takes a brief look at how the principles of this technology can still be applied in some fashion to such creative material.

25 The easiest way to apply the principles of this technology to these "noiseless" graphics is to convert them into a form which is amenable to the application of the principles of this technology. Many terms have been used in the industry for this conversion, including "ripping" a vector graphic (raster image processing) such that a vector graphic file is converted to a greyscale pixel-based raster image. Programs such as Photoshop by Adobe have such internal tools to convert vector graphic files into RGB or greyscale digital images. Once these files are in such a form, the principles of this technology can be applied in a straightforward manner. Likewise, very-low-indexed bitmaps can be converted to an RGB digital image or an equivalent. In the RGB domain, the signatures can be applied to the three color channels in appropriate ratios, or the RGB image can be simply converted into a greyscale/chroma format such as "Lab" in Adobe's Photoshop software, and the signatures can be applied to the "Lightness channel" therein. Since most of the distribution media, such as videotapes, CD-ROMs, MPEG video, digital images, and print are

- 67 -

all in forms which are amenable to the application of the principles of this technology, this conversion from vector graphic form and very-low-order graphic form is often done in any event.

Another way to apply the principles of this technology to vector graphics and very-low-order bitmapped graphics is to recognize that, indeed, there are certain properties to these inherent graphic formats which - to the eye - appear as noise. The primary example is the borders and contours between where a given line or figure is drawn or not drawn, or exactly where a bit-map changes from green to blue. In most cases, a human viewer of such graphics will be keenly aware of any attempts to "modulate signature signals" via the detailed and methodical changing of the precise contours of a graphic object. Nevertheless, such encoding of the signatures is indeed possible. The distinction between this approach and that disclosed in the bulk of this disclosure is that now the signatures must ultimately derive from what already exists in a given graphic, rather than being purely and separately created and added into a signal. This disclosure points out the possibilities here nonetheless. The basic idea is to modulate a contour, a touch right or a touch left, a touch up or a touch down, in such a way as to communicate an N-bit identification word. The locations of the changes contours would be contained in a an analogous master noise image, though now the noise would be a record of random spatial shifts one direction or another, perpendicular to a given contour. Bit values of the N-bit identification word would be encoded, and read, using the same polarity checking method between the applied change and the change recorded in the master noise image.

20 PLASTIC CREDIT AND DEBIT CARD SYSTEMS BASED ON THE PRINCIPLES OF THE TECHNOLOGY

Growth in the use of plastic credit cards, and more recently debit cards and ATM cash cards, needs little introduction. Nor does there need to be much discussion here about the long history of fraud and illicit uses of these financial instruments. The development of the credit card hologram, and its subsequent forgery development, nicely serves as a historic example of the give and take of plastic card security measures and fraudulent countermeasures. This section will concern itself with how the principles of this technology can be realized in an alternative, highly fraud-proof yet cost effective plastic card-based financial network.

A basic list of desired features for an ubiquitous plastic economy might be as follows: 1) A given plastic financial card is completely impossible to forge; 2) An attempted forged card (a "look-alike") cannot even function within a transaction setting; 3) Intercepted electronic transactions by a would-be thief would not in any way be useful or re-useable; 4) In the event of physical theft of an actual valid card, there are still formidable obstacles to a thief using that card; and 5) The overall economic cost of implementation of the financial card network is equal to or less than that of the current international credit card networks, i.e., the fully loaded cost per transaction is equal to or less than the current norm, allowing

for higher profit margins to the implementors of the networks. Apart from item 5, which would require a detailed analysis of the engineering and social issues involved with an all out implementation strategy, the following use of the principles of this technology may well achieve the above list, even item 5.

5 Figs. 22 through 26, along with the ensuing written material, collectively outline what is referred to in Fig. 26 as "The Negligible-Fraud Cash Card System." The reason that the fraud-prevention aspects of the system are highlighted in the title is that fraud, and the concomitant lost revenue therefrom, is apparently a central problem in today's plastic card based economies. The differential advantages and disadvantages of this system relative to current systems will be discussed after an illustrative embodiment is presented.

10 Fig. 22 illustrates the basic unforgeable plastic card which is quite unique to each and every user. A digital image 940 is taken of the user of the card. A computer, which is hooked into the central accounting network, 980, depicted in Fig. 26, receives the digital image 940, and after processing it (as will be described surrounding Fig. 24) produces a final rendered image which is then printed out onto the personal cash card 950. Also depicted in Fig. 22 is a straightforward identification marking, in this case a
15 bar code 952, and optional position fiducials which may assist in simplifying the scanning tolerances on the Reader 958 depicted in Fig. 23.

The short story is that the personal cash card 950 actually contains a very large amount of information unique to that particular card. There are no magnetic strips involved, though the same principles can certainly be applied to magnetic strips, such as an implanted magnetic noise signal (see
20 earlier discussion on the "fingerprinting" of magnetic strips in credit cards; here, the fingerprinting would be prominent and proactive as opposed to passive). In any event, the unique information within the image on the personal cash card 950 is stored along with the basic account information in a central accounting network, 980, Fig. 26. The basis for unbreakable security is that during transactions, the central network need only query a small fraction of the total information contained on the card, and never needs to query
25 the same precise information on any two transactions. Hundreds if not thousands or even tens of thousands of unique and secure "transaction tokens" are contained within a single personal cash card. Would-be pirates who went so far as to pick off transmissions of either encrypted or even unencrypted transactions would find the information useless thereafter. This is in marked distinction to systems which have a single complex and complete "key" (generally encrypted) which needs to be accessed, in its
30 entirety, over and over again. The personal cash card on the other hand contains thousands of separate and secure keys which can be used once, within milliseconds of time, then forever thrown away (as it were). The central network 980 keeps track of the keys and knows which have been used and which haven't.

Fig. 23 depicts what a typical point-of-sale reading device, 958, might look like. Clearly, such a
35 device would need to be manufacturable at costs well in line with, or cheaper than, current cash register

systems, ATM systems, and credit card swipers. Not depicted in Fig. 23 are the innards of the optical scanning, image processing, and data communications components, which would simply follow normal engineering design methods carrying out the functions that are to be described henceforth and are well within the capability of artisans in these fields. The reader 958 has a numeric punch pad 962 on it, showing that a normal personal identification number system can be combined with the overall design of this system adding one more conventional layer of security (generally after a theft of the physical card has occurred). It should also be pointed out that the use of the picture of the user is another strong (and increasingly common) security feature intended to thwart after-theft and illicit use. Functional elements such as the optical window, 960, are shown, mimicking the shape of the card, doubling as a centering mechanism for the scanning. Also shown is the data line cable 966 presumably connected either to a proprietor's central merchant computer system or possibly directly to the central network 980. Such a reader may also be attached directly to a cash register which performs the usual tallying of purchased items. Perhaps overkill on security would be the construction of the reader, 958, as a type of Faraday cage such that no electronic signals, such as the raw scan of the card, can emanate from the unit. The reader 958 does need to contain, preferably, digital signal processing units which will assist in swiftly calculating the dot product operations described henceforth. It also should contain local read-only memory which stores a multitude of spatial patterns (the orthogonal patterns) which will be utilized in the "recognition" steps outlined in Fig. 25 and its discussion. As related in Fig. 23, a consumer using the plastic card merely places their card on the window to pay for a transaction. A user could choose for themselves if they want to use a PIN number or not. Approval of the purchase would presumably happen within seconds, provided that the signal processing steps of Fig. 25 are properly implemented with effectively parallel digital processing hardware.

Fig. 24 takes a brief look at one way to process the raw digital image, 940, of a user into an image with more useful information content and uniqueness. It should be clearly pointed out that the raw digital image itself could in fact be used in the following methods, but that placing in additional orthogonal patterns into the image can significantly increase the overall system. (Orthogonal means that, if a given pattern is multiplied by another orthogonal pattern, the resulting number is zero, where "multiplication of patterns" is meant in the sense of vector dot products; these are all familiar terms and concepts in the art of digital image processing). Fig. 24 shows that the computer 942 can, after interrogating the raw image 970, generate a master snowy image 972 which can be added to the raw image 970 to produce a yet-more unique image which is the image that is printed onto the actual personal cash card, 950. The overall effect on the image is to "texturize" the image. In the case of a cash card, invisibility of the master snowy pattern is not as much of a requirement as with commercial imagery, and one of the only criteria for keeping the master snowy image somewhat lighter is to not obscure the image of the user. The central network, 980, stores the final processed image in the record of the account of the user, and it is this

- 70 -

unique and securely kept image which is the carrier of the highly secure "throw-away transaction keys." This image will therefore be "made available" to all duly connected point-of-sale locations in the overall network. As will be seen, none of the point-of-sale locations ever has knowledge of this image, they merely answer queries from the central network.

5 Fig. 25 steps through a typical transaction sequence. The figure is laid out via indentations, where the first column are steps performed by the point-of-sale reading device 958, the second column has information transmission steps communicated over the data line 966, and the third column has steps taken by the central network 980 which has the secured information about the user's account and the user's unique personal cash card 950. Though there is some parallelism possible in the implementation of the steps, as is normally practiced in the engineering implementation of such systems, the steps are
10 nevertheless laid out according to a general linear sequence of events.

Step one of Fig. 25 is the standard "scanning" of a personal cash card 950 within the optical window 960. This can be performed using linear optical sensors which scan the window, or via a two dimensional optical detector array such as a CCD. The resulting scan is digitized into a grey scale image
15 and stored in an image frame memory buffer such as a "framegrabber," as is now common in the designs of optical imaging systems. Once the card is scanned, a first image processing step would probably be locating the four fiducial center points, 954, and using these four points to guide all further image processing operations (i.e. the four fiducials "register" the corresponding patterns and barcodes on the personal cash card). Next, the barcode ID number would be extracted using common barcode reading
20 image processing methods. Generally, the user's account number would be determined in this step.

Step two of Fig. 25 is the optional typing in of the PIN number. Presumably most users would opt to have this feature, except those users who have a hard time remembering such things and who are convinced that no one will ever steal their cash card.

Step three of Fig. 25 entails connecting through a data line to the central accounting network and
25 doing the usual communications handshaking as is common in modem-based communications systems. A more sophisticated embodiment of this system would obviate the need for standard phone lines, such as the use of optical fiber data links, but for now we can assume it is a garden variety belltone phone line and that the reader 958 hasn't forgotten the phone number of the central network.

After basic communications are established, step four shows that the point-of-sale location
30 transmits the ID number found in step 1, along with probably an encrypted version of the PIN number (for added security, such as using the ever more ubiquitous RSA encryption methods), and appends the basic information on the merchant who operates the point-of-sale reader 958, and the amount of the requested transaction in monetary units.

Step five has the central network reading the ID number, routing the information accordingly to
35 the actual memory location of that user's account, thereafter verifying the PIN number and checking that

- 71 -

the account balance is sufficient to cover the transaction. Along the way, the central network also accesses the merchant's account, checks that it is valid, and readies it for an anticipated credit.

Step six begins with the assumption that step five passed all counts. If step five didn't, the exit step of sending a NOT OK back to the merchant is not depicted. So, if everything checks out, the central network generates twenty four sets of sixteen numbers, where all numbers are mutually exclusive, and in general, there will be a large but quite definitely finite range of numbers to choose from. Fig. 25 posits the range being 64K or 65536 numbers. It can be any practical number, actually. Thus, set one of the twenty four sets might have the numbers 23199, 54142, 11007, 2854, 61932, 32879, 38128, 48107, 65192, 522, 55723, 27833, 19284, 39970, 19307, and 41090, for example. The next set would be similarly random, but the numbers of set one would be off limits now, and so on through the twenty four sets. Thus, the central network would send (16x24x2 bytes) of numbers or 768 bytes. The actual amount of numbers can be determined by engineering optimization of security versus transmission speed issues. These random numbers are actually indexes to a set of 64K universally *a priori* defined orthogonal patterns which are well known to both the central network and are permanently stored in memory in all of the point-of-sale readers. As will be seen, a would-be thief's knowledge of these patterns is of no use.

Step seven then transmits the basic "OK to proceed" message to the reader, 958, and also sends the 24 sets of 16 random index numbers.

Step eight has the reader receiving and storing all these numbers. Then the reader, using its local microprocessor and custom designed high speed digital signal processing circuitry, steps through all twenty four sets of numbers with the intention of deriving 24 distinct floating point numbers which it will send back to the central network as a "one time key" against which the central network will check the veracity of the card's image. The reader does this by first adding together the sixteen patterns indexed by the sixteen random numbers of a given set, and then performing a common dot product operation between the resulting composite pattern and the scanned image of the card. The dot product generates a single number (which for simplicity we can call a floating point number). The reader steps through all twenty four sets in like fashion, generating a unique string of twenty four floating point numbers.

Step nine then has the reader transmitting these results back to the central network.

Step ten then has the central network performing a check on these returned twenty four numbers, presumably doing its own exact same calculations on the stored image of the card that the central network has in its own memory. The numbers sent by the reader can be "normalized," meaning that the highest absolute value of the collective twenty four dot products can be divided by itself (its unsigned value), so that brightness scale issues are removed. The resulting match between the returned values and the central network's calculated values will either be well within given tolerances if the card is valid, and way off if the card is a phony or if the card is a crude reproduction.

Step eleven then has the central network sending word whether or not the transaction was OK, and letting the customer know that they can go home with their purchased goods.

Step twelve then explicitly shows how the merchant's account is credited with the transaction amount.

5 As already stated, the primary advantage of this plastic card is to significantly reduce fraud, which apparently is a large cost to current systems. This system reduces the possibility of fraud only to those cases where the physical card is either stolen or very carefully copied. In both of these cases, there still remains the PIN security and the user picture security (a known higher security than low wage clerks analyzing signatures). Attempts to copy the card must be performed through "temporary theft" of the
10 card, and require photo-quality copying devices, not simple magnetic card swipers. The system is founded upon a modern 24 hour highly linked data network. Illicit monitoring of transactions does the monitoring party no use whether the transmissions are encrypted or not.

It will be appreciated that the foregoing approach to increasing the security of transactions involving credit and debit card systems is readily extended to any photograph-based identification system.
15 Moreover, the principles of the present technology may be applied to detect alteration of photo ID documents, and to generally enhance the confidence and security of such systems. In this regard, reference is made to Fig. 28, which depicts a photo-ID card or document 1000 which may be, for example, a passport, visa, permanent resident card ("green card"), driver's license, credit card, government employee identification, or a private industry identification badge. For convenience, such
20 photograph-based identification documents will be collectively referred to as photo ID documents.

The photo ID document includes a photograph 1010 that is attached to the document 1000. Printed, human-readable information 1012 is incorporated in the document 1000, adjacent to the photograph 1010. Machine readable information, such as that known as "bar code" may also be included adjacent to the photograph.

25 Generally, the photo ID document is constructed so that tampering with the document (for example, swapping the original photograph with another) should cause noticeable damage to the card. Nevertheless, skilled forgers are able to either alter existing documents or manufacture fraudulent photo ID documents in a manner that is extremely difficult to detect.

As noted above, the present technology enhances the security associated with the use of photo ID
30 documents by supplementing the photographic image with encoded information (which information may or may not be visually perceptible), thereby facilitating the correlation of the photographic image with other information concerning the person, such as the printed information 1012 appearing on the document 1000.

In one embodiment, the photograph 1010 may be produced from a raw digital image to which is
35 added a master snowy image as described above in connection with Figs. 22-24. The above-described

central network and point-of-sale reading device (which device, in the present embodiment, may be considered as a point-of-entry or point-of-security photo ID reading device), would essentially carry out the same processing as described with that embodiment, including the central network generation of unique numbers to serve as indices to a set of defined orthogonal patterns, the associated dot product operation
5 carried out by the reader, and the comparison with a similar operation carried out by the central network. If the numbers generated from the dot product operation carried out by the reader and the central network match, in this embodiment, the network sends the OK to the reader, indicating a legitimate or unaltered photo ID document.

In another embodiment, the photograph component 1010 of the identification document 1000 may
10 be digitized and processed so that the photographic image that is incorporated into the photo ID document 1000 corresponds to the "distributable signal" as defined above. In this instance, therefore, the photograph includes a composite, embedded code signal, imperceptible to a viewer, but carrying an N-bit identification code. It will be appreciated that the identification code can be extracted from the photo using any of the decoding techniques described above, and employing either universal or custom codes, depending upon the
15 level of security sought.

It will be appreciated that the information encoded into the photograph may correlate to, or be redundant with, the readable information 1012 appearing on the document. Accordingly, such a document could be authenticated by placing the photo ID document on a scanning system, such as would be available at a passport or visa control point. The local computer, which may be provided with the universal code
20 for extracting the identification information, displays the extracted information on the local computer screen so that the operator is able to confirm the correlation between the encoded information and the readable information 1012 carried on the document.

It will be appreciated that the information encoded with the photograph need not necessarily correlate with other information on an identification document. For example, the scanning system may
25 need only to confirm the existence of the identification code so that the user may be provided with a "go" or "no go" indication of whether the photograph has been tampered with. It will also be appreciated that the local computer, using an encrypted digital communications line, could send a packet of information to a central verification facility, which thereafter returns an encrypted "go" or "no go" indication.

In another embodiment, it is contemplated that the identification code embedded in the photograph
30 may be a robust digital image of biometric data, such as a fingerprint of the card bearer, which image, after scanning and display, may be employed for comparison with the actual fingerprint of the bearer in very high security access points where on-the-spot fingerprint recognition systems (or retinal scans, etc.) are employed.

It will be appreciated that the information embedded in the photograph need not be visually hidden
35 or steganographically embedded. For example, the photograph incorporated into the identification card

- 74 -

may be a composite of an image of the individual and one-, or two-dimensional bar codes. The bar code information would be subject to conventional optical scanning techniques (including internal cross checks) so that the information derived from the code may be compared, for example, to the information printed on the identification document.

5 It is also contemplated that the photographs of ID documents currently in use may be processed so that information correlated to the individual whose image appears in the photograph may be embedded. In this regard, the reader's attention is directed to the foregoing portion of this description entitled "Use in Printing, Paper, Documents, Plastic-Coated Identification Cards, and Other Material Where Global Embedded Codes Can Be Imprinted," wherein there is described numerous approaches to modulation of
10 physical media that may be treated as "signals" amenable to application of the present technology principles.

NETWORK LINKING METHOD USING INFORMATION

EMBEDDED IN DATA OBJECTS THAT HAVE INHERENT NOISE

15 The diagram of Fig. 27 illustrates the aspect of the technology that provides a network linking method using information embedded in data objects that have inherent noise. In one sense, this aspect is a network navigation system and, more broadly, a massively distributed indexing system that embeds addresses and indices directly within data objects themselves. As noted, this aspect is particularly well-adapted for establishing hot links with pages presented on the World Wide Web (WWW). A given data
20 object effectively contains both a graphical representation and embedded URL address.

As in previous embodiments, this embedding is carried out so that the added address information does not affect the core value of the object so far as the creator and audience are concerned. As a consequence of such embedding, only one class of data objects is present rather than the two classes (data object and discrete header file) that are attendant with traditional WWW links. The advantages of reducing
25 a hot-linked data object to a single class were mentioned above, and are elaborated upon below. In one embodiment of the technology, the World Wide Web is used as a pre-existing hot link based network. The common apparatus of this system is networked computers and computer monitors displaying the results of interactions when connected to the web. This embodiment of the technology contemplates steganographically embedding URL or other address-type information directly into images, videos, audio,
30 and other forms of data objects that are presented to a web site visitor, and which have "gray scale" or "continuous tones" or "gradations" and, consequently, inherent noise. As noted above, there are a variety of ways to realize basic steganographic implementations, all of which could be employed in accordance with the present technology.

With particular reference to Fig. 27, images, quasi-continuous tone graphics, multimedia video
35 and audio data are currently the basic building blocks of many sites 1002, 1004 on the World Wide Web.

Such data will be hereafter collectively referred to as creative data files or data objects. For illustrative purposes, a continuous-tone graphic data object 1006 (diamond ring with background) is depicted in Fig. 27.

Web site tools - both those that develop web sites 1008 and those that allow browsing them 1010 -
5 routinely deal with the various file formats in which these data objects are packaged. It is already common to distribute and disseminate these data objects 1006 as widely as possible, often with the hope on a creator's part to sell the products represented by the objects or to advertise creative services (e.g., an exemplary photograph, with an 800 phone number displayed within it, promoting a photographer's skills and service). Using the methods of this technology, individuals and organizations who create and
10 disseminate such data objects can embed an address link that leads right back to their own node on a network, their own site on the WWW.

A user at one site 1004 needs merely to point and click at the displayed object 1006. The software 1010 identifies the object as a hot link object. The software reads the URL address that is embedded within the object and routes the user to the linked web site 1002, just as if the user had used a
15 conventional web link. That linked site 1002 is the home page or network node of the creator of the object 1006, which creator may be a manufacturer. The user at the first site 1004 is then presented with, for example, an order form for purchasing the product represented by the object 1006.

It will be appreciated that the creators of objects 1006 having embedded URL addresses or indices (which objects may be referred to as "hot objects") and the manufacturers hoping to advertise their goods
20 and services can now spread their creative content like dandelion seeds in the wind across the WWW, knowing that embedded within those seeds are links back to their own home page.

It is contemplated that the object 1006 may include a visible icon 1012 (such as the exemplary "HO" abbreviation shown in Fig. 27) incorporated as part of the graphic. The icon or other subtle indicia would apprise the user that the object is a hot object, carrying the embedded URL address or other
25 information that is accessible via the software 1010.

Any human-perceptible indicium (e.g., a short musical tone) can serve the purpose of apprising the user of the hot object. It is contemplated, however, that no such indicium is required. A user's trial-and-error approach to clicking on a data object having no embedded address will merely cause the software to look for, but not find, the URL address.

30 The automation process inherent in the use of this aspect of the technology is very advantageous. Web software and web site development tools simply need to recognize this new class of embedded hot links (hot objects), operating on them in real time. Conventional hot links can be modified and supplemented simply by "uploading" a hot object into a web site repository, never requiring a web site programmer to do a thing other than basic monitoring of traffic.

- 76 -

A method of implementing the above described functions of the present technology generally involves the steps of (1) creating a set of standards by which URL addresses are steganographically embedded within images, video, audio, and other forms of data objects; and (2) designing web site development tools and web software such that they recognize this new type of data object (the hot object), the tools being designed such that when the objects are presented to a user and that user points and clicks on such an object, the user's software knows how to read or decode the steganographic information and route the user to the decoded URL address.

The foregoing portions of this description detailed a steganographic implementation (see, generally, Fig. 2 and the text associated therewith) that is readily adapted to implement the present technology. In this regard, the otherwise conventional site development tool 1008 is enhanced to include, for example, the capability to encode a bit-mapped image file with an identification code (URL address, for example) according to the present technology. In the present embodiment, it is contemplated that the commercial or transaction based hot objects may be steganographically embedded with URL addresses (or other information) using any of the universal codes described above.

The foregoing portions of this description also detailed a technique for reading or decoding steganographically embedded information (see, generally, Fig. 3 and the text associated therewith) that is readily adapted to implement the present technology. In this regard, the otherwise conventional user software 1010 is enhanced to include, for example, the capability to analyze encoded bit-mapped files and extract the identification code (URL address, for example).

While an illustrative implementation for steganographically embedding information on a data object has been described, one of ordinary skill will appreciate that any one of the multitude of available steganographic techniques may be employed to carry out the function of the present embodiment.

It will be appreciated that the present embodiment provides an immediate and common sense mechanism whereby some of the fundamental building blocks of the WWW, namely images and sound, can also become hot links to other web sites. Also, the programming of such hot objects can become fully automated merely through the distribution and availability of images and audio. No real web site programming is required. The present embodiment provides for the commercial use of the WWW in such a way that non-programmers can easily spread their message merely by creating and distributing creative content (herein, hot objects). As noted, one can also transition web based hot links themselves from a more arcane text based interface to a more natural image based interface.

Encapsulated Hot Link File Format

As noted above, once steganographic methods of hot link navigation take hold, then, as new file formats and transmission protocols develop, more traditional methods of "header-based" information attachment can enhance the basic approach built by a steganographic-based system. One way to begin

- 77 -

extending the steganographic based hot link method back into the more traditional header-based method is to define a new class of file format which could effectively become the standard class used in network navigation systems. It will be seen that objects beyond images, audio and the like can now become "hot objects", including text files, indexed graphic files, computer programs, and the like.

5 The encapsulated hot link (EHL) file format simply is a small shell placed around a large range of pre-existing file formats. The EHL header information takes only the first N bytes of a file, followed by a full and exact file in any kind of industry standard format. The EHL super-header merely attaches the correct file type, and the URL address or other index information associating that object to other nodes on a network or other databases on a network.

10 It is possible that the EHL format could be the method by which the steganographic methods are slowly replaced (but probably never completely). The slowness pays homage to the idea that file format standards often take much longer to create, implement, and get everybody to actually use (if at all). Again, the idea is that an EHL-like format and system built around it would bootstrap onto a system setup based on steganographic methods.

15

Self Extracting Web Objects

Generally speaking, three classes of data can be steganographically embedded in an object: a number (e.g. a serial or identification number, encoded in binary), an alphanumeric message (e.g. a human readable name or telephone number, encoded in ASCII or a reduced bit code), or computer instructions (e.g. JAVA or HTML instructions). The embedded URLs and the like detailed above begin to explore this third class, but a more detailed exposition of the possibilities may be helpful.

Consider a typical web page, shown in Fig. 27A. It may be viewed as including three basic components: images (#1 - #6), text, and layout.

Applicant's technology can be used to consolidate this information into a self-extracting object, and regenerate the web page from this object.

25 In accordance with this example, Fig. 27B shows the images of the Fig. 27A web page fitted together into a single RGB mosaiced image. A user can perform this operation manually using existing image processing programs, such as Adobe's Photoshop software, or the operation can be automated by a suitable software program.

30 Between certain of the image tiles in the Fig. 27B mosaic are empty areas (shown by cross-hatching).

This mosaiced image is then steganographically encoded to embed the layout instructions (e.g. HTML) and the web page text therein. In the empty areas the encoding gain can be maximized since there is no image data to corrupt. The encoded, mosaiced image is then JPEG compressed to form a self extracting web page object.

35

- 78 -

(JPEG compression is used in this example as a *lingua franca* of image representations. Another such candidate is the GIF file format. Such formats are supported by a variety of software tools and languages, making them well suited as "common carriers" of embedded information. Other image representations can of course be used.)

5 These objects can be exchanged as any other JPEG images. When the JPEG file is opened, a suitably programmed computer can detect the presence of the embedded information and extract the layout data and text. Among other information, the layout data specifies where the images forming the mosaic are to be located in the final web page. The computer can follow the embedded HTML instructions to recreate the original web page, complete with graphics, text, and links to other URLs.

10 If the self extracting web page object is viewed by a conventional JPEG viewer, it does not self-extract. However, the user will see the logos and artwork associated with the web page (with noise-like "grouting" between certain of the images). Artisans will recognize that this is in stark contrast to viewing of other compressed data objects (e.g. PKZIP files and self extracting text archives) which typically appear totally unintelligible unless fully extracted.

15 (The foregoing advantages can largely be achieved by placing the web page text and layout instructions in a header file associated with a JPEG-compressed mosaiced image file. However, industry standardization of the header formats needed to make such systems practical appears difficult, if not impossible.)

20 Palettes of Steganographically Encoded Images

 Once web images with embedded URL information become widespread, such web images can be collected into "palettes" and presented to users as high level navigation tools. Navigation is effected by clicking on such images (e.g. logos for different web pages) rather than clicking on textual web page names. A suitably programmed computer can decode the embedded URL information from the selected
25 image, and establish the requested connection.

 In one embodiment, self-extraction of the above-described web page objects automatically generates thumbnail images corresponding to the extracted pages (e.g. representative logos), which are then stored in a subdirectory in the computer's file system dedicated to collecting such thumbnails. In each such thumbnail is embedded a URL, such as the URL of the extracted page or the URL of the site from
30 which the self-extracting object was obtained. This subdirectory can be accessed to display a palette of navigational thumbnails for selection by the user.

Specific Example of a Computer System Linked by Using Information in Data Objects

 The present invention allows a digital watermark to be embedded directly into photographs, video,
35 computer images, audio, and other forms of creative property. The range of applications is amazingly

- 79 -

wide, and includes things like video and movie content watermarking for proof of ownership and tracking, watermarking music and audio content, even texturing physical materials such as auto parts. As discussed above, this watermark is imperceptible to the eye, and imperceptible to the ear, but a computer analysis can read the watermark and discover a message it carries. The watermark is repeated throughout the
5 property. It is robust, and typically survives multiple generations of copying, modification, printing, scanning, and compression. A watermark may carry a copyright notice, a unique serial number, a Transaction ID, as well as other application specific data.

Conceptually, the example below is made with respect to a watermark embedded by "PictureMarc" which illustrates the idea by carrying a unique creator id, and a set of image
10 attributes. This creator id corresponds to complete contact information available through an on-line service, MarcCentre. Technically, the watermark is 128 bits long and has a data payload of 76 bits, although the invention is not restricted to any particular watermark size. In this example, it is short so that it can be repeated many times through an image, making it possible to find the message even from a portion of the original image. Other approaches can include data such as an
15 image serial number or a Transaction ID in the watermark message. It is also possible to add a person's name or other such data (to a computer, it is all data), but one application is focused on embedding a more structured message. First, to accommodate a text string, there would be a tendency to push for longer messages. Long messages result in not being able to fit very many copies of the message in an image, and thus the watermark becomes less robust. Also, just
20 putting a name in an image is much like the credit line in a magazine: A. Jones tells who the person is, but not how to contact the person. With a watermark system available on line, a user can get complete and up to date contact details from each Creator ID through the on-line service. Accordingly, it is preferable to associate persistent digital information with an image. The core software has been designed to enable working with vendors or large customers who may want to
25 develop specific watermark message types for special applications such as embedding the film speed and exposure in the image, or serializing images in a large collection for internal tracking and management.

One approach to implementing this embodiment involves "PictureMarc", which includes a writer portion to embed watermarks and a reader portion to find them and read their contents.
30 This bundle is integrated into Adobe Photoshop® 4.0 (Adobe Photoshop is a trademark of Adobe Systems, Inc.) as a plugin extension, and from the end-user's perspective is free. A free reader will also be available through a web site. PictureMarc for Photoshop is supported on Apple

- 80 -

Macintosh and Microsoft Windows platforms. PictureMarc supports embedding a Creator ID and image attributes in an image. PictureMarc also has the ability to carry a unique Image ID or even a Transaction ID in the watermark. These formats will support tracking images both to a collection, and to the person and terms that were in place when the image sold. The watermark
5 reader will accept and read all watermark formats.

As an example, the first step is to read an image from a CD-ROM, download a file to a local computer over the Internet or convert a hardcopy image, such as a photograph, into digital by scanning the image with any of the common image scanners.. Then start an image manipulation program that supports watermarking, and load the image file from a file selection
10 menu. Select the menu item to add a watermark, and save the watermarked image into a file.

To recover a watermarked image from a digital image on, say, a PC, just as in the case of embedding a watermark, convert the image to a suitable computer file format, scanning it in if needed. Then start either an image manipulation tool supporting watermarking, or start a stand-alone reader. Now open the image. In the case of a preferred embodiment image manipulation
15 tool, the tool will automatically check for a watermark when the digital image file is opened, copied to the system clipboard or the image is scanned into the computer. If present, a copyright symbol is added to the title or status bar, or some other signal is provided to the user. By selecting a read watermark menu option, a user can discover the contents of the watermark. From the results dialog, a user can click on the Web Lookup button to retrieve complete contact
20 information from MarcCentre, a network WWW on-line service. The application can be carried out with at least two kinds of watermarks. The first is a public watermark, which can be read by any suitably configured reader. This is the type of watermark supported by PictureMarc--Its purpose is communication. The second is a private watermark. Private watermarks are associated with a secret key used at the time the watermark is embedded in an image. Only the person with
25 the secret key can find and read a private watermark. Private watermarks are used in proof of ownership applications. Both a public and a private watermark can coexist in an image.

If a user attempts to read a watermark and none is there, the reader portion will simply report that a watermark was not found in the image. If a user attempts to watermark an image already watermarked, the user will be informed there is already a watermark in the image. and

- 81 -

will be prevented from adding another watermark. In creating a composite image, the user is first told about the presence of a watermark in each watermarked original (assuming the use of a tool supporting watermarks). This addresses the first order problem of communicating copyright. The watermark will not prevent creating a composite work, in that a watermarked image, or a portion thereof, can be used as an element of a composite image. However, if the composite image contains multiple watermarked elements, the reader portion will tend to read the watermark from the most dominant piece. However, if the user suspects a component to be watermarked, or is simply interested in the source of an image element, the user can check for the watermark in that element by highlighting the area in question, and invoking the read watermark function. This provides a tool to inspect work to discover sources of individual elements. Copyright hierarchies in composited works and partnering with companies addressing rights management can all be handled in such a network application of a watermark.

Preferably, applications which include watermarking, such as Photoshop 4.0, support automatic detection ("autodetect"). This means that whenever an image is opened or scanned into the application, a quick check is done to determine if a watermark is present. If it is present, a visual indication, such as a copyright symbol, is added to the title or status bar.

The watermark placed by PictureMarc is a 128 bits long, with a data payload of 76 bits. The remaining bits are dedicated to header and control information, and error detection and correction. However, the bandwidth can be increased and consequently the data payload can also be increased. The detect process looks for the presence of a watermark, which completes in a couple of seconds, usually under 3 seconds regardless of image size, and the typical time to perform a full read watermark operation is under 15 seconds.

PictureMarc supports all the file formats of the host application, since it works with pixels in each of the major color schemes (CMYK, RGB, LAB, Grayscale). In the case of PictureMarc for Photoshop, all of the Photoshop file formats are supported.

Note that the watermark itself carries version and message type information, much like a communications message. PictureMarc has been designed to support upward compatibility. Any

- 82 -

suitably configured reader will be able to read previous versions of the watermark. If a reader finds a newer version watermark which it does not support, the user will be notified and told to download the most recent reader.

In the present application of PictureMarc, an image generally needs to be at least 256x256 pixels to carry a watermark. Often a watermark can be found in patches as small as 100 pixels square. Arbitrarily large images can be watermarked, constrained only by available memory and disk space. The watermark does not add data to the image file, since it only makes very small changes to the luminance of pixels. This can change the size of a compressed image by a few percent, but the size of uncompressed images, such as bitmaps (.BMPs), are unchanged. While a watermark survives compression, to watermark an already compressed image, it is preferable to either work from the original, or uncompress, watermark, and recompress the image. Future versions will likely support adding watermarks directly to the JPEG compressed file.

PictureMarc supports the automation features provided by Adobe Photoshop 4.0. This means that repetitive watermarking operations can be automated now. In other applications, support can be provided for unattended batch processing of images, aimed to support the needs of image distributors, especially those companies looking to watermark images as they are distributed with unique Transition IDs.

As to the use of a network on line service, an image creator or distributor subscribes to the MarcCentre, on-line locator service. Subscribing includes providing a set of contact information to be made available when one of a user's watermarked images is found. Contact details include name, address, phone number, specialty, intermediary (e.g., organization, stock agency, representative, etc.) representing the creator's work and so on. When subscribing, the user is given a unique Creator ID. The user provides this Creator ID to a copy of PictureMarc, one time when first using the software-based system. This Creator ID links all of the user's images to facilitate contact details. Information can be queried via the Internet, or via a fax-back service. Querying is preferably free.

A subscription to MarcCentre is priced to include unlimited access to contact information by anyone finding watermarked images. A Creator ID in the image leverages the power of a network accessible database, and lets the image creator communicate much more information than could fit in the image. In essence, we are creating a copyright communication system.

- 5 Watermark detection in programs such as Photoshop is automatic. The Creator ID in the image, along with the image attributes, uniquely identify the image creator or distributor. The contact information corresponding to the Creator ID is always up to date, no matter how long ago a watermarked image was distributed. To make this copyright communication system work, we need unique ids which correspond to the appropriate image creator contact information.

- 10 If there is no network connection for PictureMarc, an equivalent approach can be used: contact information can be communicated via a fax-back service. This can be carried out by punching in the Creator ID on a touch-tone keypad with telephony, and the resulting contact details for that image creator will be sent by fax.

Contact details will be stored on a robust server being hosted by an Internet provider.

- 15 The Web site is responsible for maintaining the locator service and ensuring that information is available. Individual image creators are responsible for maintaining their own contact details, keeping them current as information changes. This ensures that all of the images distributed point to up-to-date information about the creator and the creator's work.

- 20 Data in the Web site repository is stored in a way that when a user enter contact details, the user has control over which pieces of information are viewable by viewers. Thus, for example, the user can designate whether to exclude his or her address, phone number, and other specific information.

- 25 It is helpful for users to receive regular reports, generated by the Web page site software, about the number of people querying a specific user's information. These can be sent to the user via e-mail or fax, to help communicate the activity concerning watermarked images.

Before turning to a specific embodiment of the present invention, it is important to stress the general applicability of embedding a watermark, then reading it to obtain data, and then using the data to link to a database--preferably via a computer network. This is applicable to all media for which this technology is useful. This includes watermarking a photograph or a series of
5 watermarks in a series of images (e.g., television, video), in sound or other analog signals (e.g., recordings, cellular telephone or other broadcasts), and other applications indicated herein. In any of these applications, while the watermark can contain self-identifying data, a centralized repository of watermark data facilitates broad, efficient usage.

Thus, the following specific description of a preferred embodiment of a central repository,
10 accessible by a network, is intended to be illustrative of the same kind of an approach for other applications and media suitable for watermarking. And it is to be understood that while the following discussion is made in the context of image processing, the same general approach can be used for handling watermarks in acoustic or other applications. Also, while the Adobe Photoshop application is used to illustrate the generally applicable concepts, it should be
15 appreciated that the present invention can be used in connection with other such products from other vendors, as well as in stand alone form.

Fig. 43 illustrates an application of the invention using such an application as Adobe Photoshop 4.0, equipped with a plug in version of the digital watermarking technology discussed herein and illustrated as "PictureMarc," running on a digital electrical computer. The
20 computer running the Adobe application is adapted for communication over a network (such as the WWW) with another digital electrical computer, running the MarcCentre Web Page computer program.

At step 1 of Fig. 43 a creator of a digitized photographic image, the "user" of the Adobe Photoshop, commences by obtaining a Creator ID. Such data is created and stored in a
25 central repository designated in Fig. 43 as MarcCentre Locator Service, which is connectable to PictureMarc via Internet or fax. At step 2 of Fig. 43, the user creates and embeds a watermark in the digital image. At step 3 of Fig. 43, that image can be stored and/or printed, possibly subsequent to distribution over a computer system such as the Internet or WWW. When the image bearing the watermark is used in digital form, at step 4 of Fig. 43, that image is examined

- 85 -

for a watermark by a computer running the PictureMarc computer program. Preferably, the examination is conducted automatically (rather than selectively) upon opening an image, e.g., upon locating the image on a computer screen clipboard. If a watermark is detected, at step 5 of Fig. 43, PictureMarc communicates with the MarcCentre Locator Service to look up creator information and present a WWW page to the user.

The means by which a user obtains a Creator ID in step 1 of Fig. 43 is elaborated in Figs. 44, 45, and 46, which represent screen images of MarcCentre, a Web Page locatable on the WWW. Introductory information is set forth in Fig. 44. Note the mention of a single mouse click. This feature refers to a hot button, such as a copyright symbol in a title bar of a screen, which permits someone using the Web Page to obtain information about the image that was pre-specified by the creator. (A hot button (pointer to information stored in a MarcCentre database, which can in turn have a pointer to another database, such as that accessed by means of the network or WWW.)

Figs. 45 and 46 show sequential screens of a "form" to be completed by computer entries. This information is stored in a MarcCentre database. Note too that the information entered to complete the form include a fee, and the fee account for the user is also stored in a MarcCentre database. After the form has been completed and stored and the fee has been secured, a Creator ID is issued by the MarcCentre service.

The means by which a user with a Creator ID embeds a registered watermark in step 2 of Fig. 43 is elaborated in Figs. 47, 48, and 49. The Photoshop software uses a filter to call up the "Digimarc" PictureMarc software and select options to embed a watermark. Bolded bars on the screen of Fig. 47 illustrate a selection being made for a digital image in the background. Other information on the screen is typical for a windows application, the Adobe Photoshop application, and or generally known information about digital images, e.g., 100% shows the scale of the image, "RGB" indicates the image format, etc.

Fig. 48 illustrates a "Dialog Box" used in associating information in connection with the embedding of the watermark. In a portion of the screen, the Copyright Information is

- 86 -

specified, including the Creator ID, the Type of Use (either Restricted or Royalty Free), Adult Content, and a Personalize option. A demonstration ID for use by those who do not have a registered Creator ID is "PictureMarc Demo." The Adult Content feature permits embedding in the watermark data, preferably content-indicative data. Here the data is used for identifying that the photograph contains adult content, and thus a subsequent use of the photograph can be restricted. However, an alternative is to include other data or other content-identifying data, such as cataloging information, in the watermark itself. Cataloging information permits subject matter searching in a collection of photographs. Another feature shown in Fig. 48 is a slidable scale for adjusting the intensity of the watermark. The more visible the watermark, the more durable the watermark is. Watermark durability (visibility) is also indicated by a number shown on the screen. Once the screen has been completed, if it is "OK", the user can click on this button and thereby embed the watermark. Otherwise, the user can select to cancel and thereby abort the process that would have led to the watermarking.

Selecting the Personalize option of Fig. 48 calls the screen illustrated in Fig. 49, which is used to obtain a personalized, registered Creator ID. After the Creator ID is entered, selection of the Register option launches a WWW browser directed by the URL address shown entered on the screen to the MarcCentre registration page.

In a preferred embodiment of step 4 of Fig. 43, each time an image is opened, scanned in, or copied to the clipboard by any application, the image is automatically tested for a watermark, as illustrated in Fig. 50. This "autodetect" approach generates a signal indicative of the result of the test, for example, one or more copyright symbols shown in the title page of the screen in Fig. 50. Automatic checking for a watermark takes less computer time than the alternative of selectively reading the watermark, which is illustrated in Figs. 51 and 52. However, the selective approach is less comprehensive in monitoring the use of watermarked images, which is disadvantageous if the application is configured to automatically report the detection back to the MarcCentre. In Fig. 50, the copyright symbol in the title area of the photograph is the previously mentioned hot button. The copyright symbol in the lower left of the screen can alternatively or also be a hot button.

Figs. 51 and 52 elaborate the alternative approach to step 4 of Fig. 43 in which an image is opened and examined for a watermark. Preferably, each time an image is opened, there is an automatic check for a watermark. In any case, Fig. 51 is similar to the screen of Fig. 47, except that a selection is shown to read for a watermark.

5 Fig. 52 shows a screen that is be produced in response to the reading a watermark. This screen provides the Creator ID of the owner and / or distributor of this image, as well as whether use of this image is restricted or royalty free. Note the Web Lookup button, which can be selected to launch a WWW browser directed by the URL address shown entered on the screen to the MarcCentre Web Page, and obtains the data in the corresponding MarcCentre
10 database corresponding to the Creator ID.

Fig. 53 illustrates a screen produced from selecting the Web Lookup feature of Fig. 53. This screen provides the contact information for the image creator, including phone number, E-mail address and Web links. Note the selectable buttons at the bottom of the screen illustrated in Fig. 53, including buttons that take the user to promotional offers, allow the user to
15 download additional software, subscribe to the MarcCentre service, update contact information, search for a creator or provide feedback on the MarcCentre service. Selecting the Creator Search option leads to the screen shown in Fig. 54. This self-explanatory screen enables a user to obtain contact information about a particular creator by entering the Creator ID.

In a variation on the above-described approach, extended information is embedded
20 in the watermark and subsequently detected, as respectively illustrated in Figs. 55 and 56. Such a feature would be inserted as an additional filter in Fig. 47. In Fig. 55., the dialog box permits inputting an Organization ID, an Item ID, and a Transaction ID. An Organization ID identifies an organization or other intermediary representing the creator and/or authorized for distributing the photograph and licensing rights thereto, such as a photo stock house. An Item ID identifies
25 the photograph in a collection, such as a CD ROM. A Transaction ID identifies the specific transaction that is authorized (e.g., publication one time only in a newspaper), which permits detection of an authorized usage. The other feature on the screen, The Bump Size specification, pertains to the smallest unit of data, in the horizontal and vertical directions, that the watermark algorithms operate on. The extended information embedded in the watermark can be detected in

- 88 -

a manner analogous to that described with respect to Fig. 48, except that the additional information is shown as illustrated in Fig. 56.

An alternative to a "plug in" embodiment is a stand alone version, which is exemplified in the next sequence of figures. As discussed above, for any application that involves a watermark, preferably there is an autodetect capability, and Fig. 57 illustrates a screen from a stand alone image reader with a clip board triggered by merely putting an image on the clipboard. In the Clipboard Viewer window of Fig. 57, a watermark detected on the image can be used to signal the user by one of several methods that can be used individually or in a combination. In one method, a title bar icon or some other bar or icon alternates between "normal" and "watermark detected" states periodically (for example, at one second intervals). Another approach is to use other state changes, like having the title bar or a task bar alternate between highlighted and normal states periodically. Still another method is to have the detection initiate a special dialog box display, as illustrated in Fig. 58. The Read button in the special dialog box (triggered by a click of a mouse, like other buttons) causes reading and display of the watermark information in the image. The Close button ends the special dialog. (So as not to interfere with image processing, it can be convenient to show the special dialog box for a limited amount of time, and a small window therein can show the time remaining before the special dialog box disappears.) Buttons shown on the screen include one resembling a file folder to open a file; the next button is for copying a selected image from a first application to the clipboard; the next button is for pasting an image from the clipboard to a second application; and the last is for help menus.

Fig. 59 shows the image after pasting it from the clipboard, in the course of image processing. The Digimarc Reader - Image 1 window pops up on the screen to permit clicking the Web lookup button, which launches the user's web browser to result in a display of the creator's web page on the MarcCentre Web site, with the rest following as above.

Note again that the foregoing approach is representative, and it can be modified for another media in which a watermark is used. For example, instead of an image, a sound can be used.

POTENTIAL USE OF THE TECHNOLOGY IN THE PROTECTION AND CONTROL OF
SOFTWARE PROGRAMS

5 The illicit use, copying, and reselling of software programs represents a huge loss of revenues to the software industry at large. The prior art methods for attempting to mitigate this problem are very broad and will not be discussed here. What will be discussed is how the principles of this technology might be brought to bear on this huge problem. It is entirely unclear whether the tools provided by this technology will have any economic advantage (all things considered) over the existing countermeasures both in place and contemplated.

10 The state of technology over the last decade or more has made it a general necessity to deliver a full and complete copy of a software program in order for that program to function on a user's computer. In effect, \$X were invested in creating a software program where X is large, and the entire fruits of that development must be delivered in its entirety to a user in order for that user to gain value from the software product. Fortunately this is generally compiled code,
15 but the point is that this is a shaky distribution situation looked at in the abstract. The most mundane (and harmless in the minds of most perpetrators) illicit copying and use of the program can be performed rather easily.

This disclosure offers, at first, an abstract approach which may or may not prove to be economical in the broadest sense (where the recovered revenue to cost ratio would exceed that of
20 most competing methods, for example). The approach expands upon the methods and approaches already laid out in the section on plastic credit and debit cards. The abstract concept begins by positing a "large set of unique patterns," unique among themselves, unique to a given product, and unique to a given purchaser of that product. This set of patterns effectively contains thousands and even millions of absolutely unique "secret keys" to use the cryptology vernacular.
25 Importantly and distinctly, these keys are non-deterministic, that is, they do not arise from singular sub-1000 or sub-2000 bit keys such as with the RSA key based systems. This large set of patterns is measured in kilobytes and Megabytes, and as mentioned, is non-deterministic in

- 90 -

nature. Furthermore, still at the most abstract level, these patterns are fully capable of being encrypted via standard techniques and analyzed within the encrypted domain, where the analysis is made on only a small portion of the large set of patterns, and that even in the worst case scenario where a would-be pirate is monitoring the step-by-step microcode instructions of a microprocessor, this gathered information would provide no useful information to the would-be pirate. This latter point is an important one when it comes to "implementation security" as opposed to "innate security" as will be briefly discussed below.

So what could be the differential properties of this type of key based system as opposed to, for example, the RSA cryptology methods which are already well respected, relatively simple, etc. etc? As mentioned earlier, this discussion is not going to attempt a commercial side-by-side analysis. Instead, we'll just focus on the differing properties. The main distinguishing features fall out in the implementation realm (the implementation security). One example is that in single low-bit-number private key systems, the mere local use and re-use of a single private key is an inherently weak link in an encrypted transmission system. ["Encrypted transmission systems" are discussed here in the sense that securing the paid-for use of software programs will in this discussion require de facto encrypted communication between a user of the software and the "bank" which allows the user to use the program; it is encryption in the service of electronic financial transactions looked at in another light.] Would-be hackers wishing to defeat so-called secure systems never attack the fundamental hard-wired security (the innate security) of the pristine usage of the methods, they attack the implementation of those methods, centered around human nature and human oversights. It is here, still in the abstract, that the creation of a much larger key base, which is itself non-deterministic in nature, and which is more geared toward effectively throw-away keys, begins to "idiot proof" the more historically vulnerable implementation of a given secure system. The huge set of keys is not even comprehensible to the average holder of those keys, and their use of those keys (i.e., the "implementation" of those keys) can randomly select keys, easily throw them out after a time, and can use them in a way that no "eavesdropper" will gain any useful information in the eavesdropping, especially when well within a millionth of the amount of time that an eavesdropper could "decipher" a key, its usefulness in the system would be long past.

Turning the abstract to the semi-concrete, one possible new approach to securely delivering a software product to ONLY the bonafide purchasers of that product is the following. In a mass economic sense, this new method is entirely founded upon a modest rate realtime digital connectivity (often, but not necessarily standard encrypted) between a user's computer
5 network and the selling company's network. At first glance this smells like trouble to any good marketing person, and indeed, this may throw the baby out with the bathwater if by trying to recover lost revenues, you lose more legitimate revenue along the way (all part of the bottom line analysis). This new method dictates that a company selling a piece of software supplies to anyone who is willing to take it about 99.8% of its functional software for local storage on a
10 user's network (for speed and minimizing transmission needs). This "free core program" is entirely unfunctional and designed so that even the craftiest hackers can't make use of it or "decompile it" in some sense. Legitimate activation and use of this program is performed purely on a instruction-cycle-count basis and purely in a simple very low overhead communications basis between the user's network and the company's network. A customer who wishes to use the
15 product sends payment to the company via any of the dozens of good ways to do so. The customer is sent, via common shipment methods, or via commonly secured encrypted data channels, their "huge set of unique secret keys." If we were to look at this large set as if it were an image, it would look just like the snowy images discussed over and over again in other parts of this disclosure. (Here, the "signature" is the image, rather than being imperceptibly placed
20 onto another image). The special nature of this large set is that it is what we might call "ridiculously unique" and contains a large number of secret keys. (The "ridiculous" comes from the simple math on the number of combinations that are possible with, say 1 Megabyte of random bit values, equaling exactly the number that "all ones" would give, thus 1 Megabyte being approximately 10 raised to the ~2,400,000 power, plenty of room for many people having many
25 throwaway secret keys). It is important to re-emphasize that the purchased entity is literally: productive use of the tool. The marketing of this would need to be very liberal in its allotment of this productivity, since per-use payment schemes notoriously turn off users and can lower overall revenues significantly.

This large set of secret keys is itself encrypted using standard encryption techniques. The
30 basis for relatively higher "implementation security" can now begin to manifest itself. Assume that the user now wishes to use the software product. They fire up the free core, and the free

- 92 -

core program finds that the user has installed their large set of unique encrypted keys. The core program calls the company network and does the usual handshaking. The company network, knowing the large set of keys belonging to that bonafide user, sends out a query on some simple set of patterns, almost exactly the same way as described in the section on the debit and credit cards. The query is such a small set of the whole, that the inner workings of the core program do not even need to decrypt the whole set of keys, only certain parts of the keys, thus no decrypted version of the keys ever exist, even within the machine cycles on the local computer itself. As can be seen, this does not require the "signatures within a picture" methods of the main disclosure, instead, the many unique keys ARE the picture. The core program interrogates the keys by performing certain dot products, then sends the dot products back to the company's network for verification. See Fig. 25 and the accompanying discussion for typical details on a verification transaction. Generally encrypted verification is sent, and the core program now "enables" itself to perform a certain amount of instructions, for example, allowing 100,000 characters being typed into a word processing program (before another unique key needs to be transmitted to enable another 100,000). In this example, a purchaser may have bought the number of instructions which are typically used within a one year period by a single user of the word processor program. The purchaser of this product now has no incentive to "copy" the program and give it to their friends and relatives.

All of the above is well and fine except for two simple problems. The first problem can be called "the cloning problem" and the second "the big brother problem." The solutions to these two problems are intimately linked. The latter problem will ultimately become a purely social problem, with certain technical solutions as mere tools not ends.

The cloning problem is the following. It generally applies to a more sophisticated pirate of software rather than the currently common "friend gives their distribution CD to a friend" kind of piracy. Crafty-hacker "A" knows that if she performs a system-state clone of the "enabled" program in its entirety and installs this clone on another machine, then this second machine effectively doubles the value received for the same money. Keeping this clone in digital storage, hacker "A" only needs to recall it and reinstall the clone after the first period is run out, thus indefinitely using the program for a single payment, or she can give the clone to their hacker friend "B" for a six-pack of beer. One good solution to this problem requires, again, a rather well

- 93 -

developed and low cost real time digital connectivity between user site and company enabling network. This ubiquitous connectivity generally does not exist today but is fast growing through the Internet and the basic growth in digital bandwidth. Part and parcel of the "enabling" is a negligible communications cost random auditing function wherein the functioning program

5 routinely and irregularly performs handshakes and verifications with the company network. It does so, on average, during a cycle which includes a rather small amount of productivity cycles of the program. The resulting average productivity cycle is in general much less than the raw total cost of the cloning process of the overall enabled program. Thus, even if an enabled program is cloned, the usefulness of that instantaneous clone is highly limited, and it would be

10 much more cost effective to pay the asking price of the selling company than to repeat the cloning process on such short time periods. Hackers could break this system for fun, but certainly not for profit. The flip side to this arrangement is that if a program "calls up" the company's network for a random audit, the allotted productivity count for that user on that program is accounted for, and that in cases where bonafide payment has not been received, the

15 company network simply withholds its verification and the program no longer functions. We're back to where users have no incentive to "give this away" to friends unless it is an explicit gift (which probably is quite appropriate if they have indeed paid for it: "do anything you like with it, you paid for it").

The second problem of "big brother" and the intuitively mysterious "enabling"

20 communications between a user's network and a company's network would as mentioned be a social and perceptual problem that should have all manner of potential real and imagined solutions. Even with the best and objectively unbeatable anti-big-brother solutions, there will still be a hard-core conspiracy theory crowd claiming it just ain't so. With this in mind, one potential solution is to set up a single program registry which is largely a public or non-profit institution to

25 handling and coordinating the realtime verification networks. Such an entity would then have company clients as well as user clients. An organization such as the Software Publishers Association, for example, may choose to lead such an effort.

Concluding this section, it should be re-emphasized that the methods here outlined require a highly connected distributed system, in other words, a more ubiquitous and inexpensive Internet

30 than exists in mid 1995. Simple trend extrapolation would argue that this is not too far off from

- 94 -

1995. The growth rate in raw digital communications bandwidth also argues that the above system might be more practical, sooner, than it might first appear. (The prospect of interactive TV brings with it the promise of a fast network linking millions of sites around the world.)

5 USE OF CURRENT CRYPTOLOGY METHODS IN CONJUNCTION WITH THIS TECHNOLOGY

It should be briefly noted that certain implementations of the principles of this technology probably can make good use of current cryptographic technologies. One case in point might be a system whereby graphic artists and digital photographers perform realtime registration of their
10 photographs with the copyright office. It might be advantageous to send the master code signals, or some representative portion thereof, directly to a third party registry. In this case, a photographer would want to know that their codes were being transmitted securely and not stolen along the way. In this case, certain common cryptographic transmission might be employed. Also, photographers or musicians, or any users of this technology, may want to have reliable time
15 stamping services which are becoming more common. Such a service could be advantageously used in conjunction with the principles of this technology.

DETAILS ON THE LEGITIMATE AND ILLEGITIMATE DETECTION AND REMOVAL OF INVISIBLE SIGNATURES

20 In general, if a given entity can recognize the signatures hidden within a given set of empirical data, that same entity can take steps to remove those signatures. In practice, the degree of difficulty between the former condition and the latter condition can be made quite large, fortunately. On one extreme, one could posit a software program which is generally very difficult to "decompile" and which does recognition functions on empirical data. This same bit of
25 software could not generally be used to "strip" the signatures (without going to extreme lengths).

- 95 -

On the other hand, if a hacker goes to the trouble of discovering and understanding the "public codes" used within some system of data interchange, and that hacker knows how to recognize the signatures, it would not be a large step for that hacker to read in a given set of signed data and create a data set with the signatures effectively removed. In this latter example, interestingly
5 enough, there would often be telltale statistics that signatures had been removed, statistics which will not be discussed here.

These and other such attempts to remove the signatures we can refer to as illicit attempts. Current and past evolution of the copyright laws have generally targeted such activity as coming under criminal activity and have usually placed such language, along with penalties and
10 enforcement language, into the standing laws. Presumably any and all practitioners of this signature technology will go to lengths to make sure that the same kind of a) creation, b) distribution, and c) use of these kinds of illicit removal of copyright protection mechanisms are criminal offenses subject to enforcement and penalty. On the other hand, it is an object of this technology to point out that through the recognition steps outlined in this disclosure, software
15 programs can be made such that the recognition of signatures can simply lead to their removal by inverting the known signatures by the amount equal to their found signal energy in the recognition process (i.e., remove the size of the given code signal by exact amount found). By pointing this out in this disclosure, it is clear that such software or hardware which performs this signature removal operation will not only (presumably) be criminal, but it will also be liable to
20 infringement to the extent that it is not properly licensed by the holders of the (presumably) patented technology.

The case of legitimate and normal recognition of the signatures is straightforward. In one example, the public signatures could deliberately be made marginally visible (i.e. their intensity would be deliberately high), and in this way a form of sending out "proof comps" can be
25 accomplished. "Comps" and "proofs" have been used in the photographic industry for quite some time, where a degraded image is purposely sent out to prospective customers so that they might evaluate it but not be able to use it in a commercially meaningful way. In the case of this technology, increasing the intensity of the public codes can serve as a way to "degrade" the commercial value intentionally, then through hardware or software activated by paying a purchase

price for the material, the public signatures can be removed (and possibly replaced by a new invisible tracking code or signature, public and/or private.

MONITORING STATIONS AND MONITORING SERVICES

5 Ubiquitous and cost effective recognition of signatures is a central issue to the broadest proliferation of the principles of this technology. Several sections of this disclosure deal with this topic in various ways. This section focuses on the idea that entities such as monitoring nodes, monitoring stations, and monitoring agencies can be created as part of a systematic enforcement of the principles of the technology. In order for such entities to operate, they require knowledge
10 of the master codes, and they may require access to empirical data in its raw (unencrypted and untransformed) form. (Having access to original unsigned empirical data helps in finer analyses but is not necessary.)

 Three basic forms of monitoring stations fall out directly from the admittedly arbitrarily defined classes of master codes: a private monitoring station, a semi-public, and a public. The
15 distinctions are simply based on the knowledge of the master codes. An example of the fully private monitoring station might be a large photographic stock house which decides to place certain basic patterns into its distributed material which it knows that a truly crafty pirate could decipher and remove, but it thinks this likelihood is ridiculously small on an economic scale. - This stock house hires a part-time person to come in and randomly check high value ads and
20 other photography in the public domain to search for these relatively easy to find base patterns, as well as checking photographs that stock house staff members have "spotted" and think it might be infringement material. The part time person cranks through a large stack of these potential infringement cases in a few hours, and where the base patterns are found, now a more thorough analysis takes place to locate the original image and go through the full process of unique
25 identification as outlined in this disclosure. Two core economic values accrue to the stock house in doing this, values which by definition will outweigh the costs of the monitoring service and the cost of the signing process itself. The first value is in letting their customers and the world know

- 97 -

that they are signing their material and that the monitoring service is in place, backed up by whatever statistics on the ability to catch infringers. This is the deterrent value, which probably will be the largest value eventually. A general pre-requisite to this first value is the actual recovered royalties derived from the monitoring effort and its building of a track record for being formidable (enhancing the first value).

The semi-public monitoring stations and the public monitoring stations largely follow the same pattern, although in these systems it is possible to actually set up third party services which are given knowledge of the master codes by clients, and the services merely fish through thousands and millions of "creative property" hunting for the codes and reporting the results to the clients. ASCAP and BMI have "lower tech" approaches to this basic service.

A large coordinated monitoring service using the principles of this technology would classify its creative property supplier clients into two basic categories, those that provide master codes themselves and wish the codes to remain secure and unpublished, and those that use generally public domain master codes (and hybrids of the two, of course). The monitoring service would perform daily samplings (checks) of publicly available imagery, video, audio, etc., doing high level pattern checks with a bank of supercomputers. Magazine ads and images would be scanned in for analysis, video grabbed off of commercial channels would be digitized, audio would be sampled, public Internet sites randomly downloaded, etc. These basic data streams would then be fed into an ever-churning monitoring program which randomly looks for pattern matches between its large bank of public and private codes, and the data material it is checking. A small sub-set, which itself will probably be a large set, will be flagged as potential match candidates, and these will be fed into a more refined checking system which begins to attempt to identify which exact signatures may be present and to perform a more fine analysis on the given flagged material. Presumably a small set would then fall out as flagged match material, owners of that material would be positively identified and a monitoring report would be sent to the client so that they can verify that it was a legitimate sale of their material. The same two values of the private monitoring service outlined above apply in this case as well. The monitoring service could also serve as a formal bully in cases of a found and proven infringement, sending out letters to infringing parties witnessing the found infringement and seeking inflated royalties so that the infringing party might avoid the more costly alternative of going to court.

METHOD FOR EMBEDDING SUBLIMINAL REGISTRATION PATTERNS INTO IMAGES
AND OTHER SIGNALS

The very notion of reading embedded signatures involves the concept of registration. The underlying master noise signal must be known, and its relative position needs to be ascertained
5 (registered) in order to initiate the reading process itself (e.g. the reading of the 1's and 0's of the N-bit identification word). When one has access to the original or a thumbnail of the unsigned signal, this registration process is quite straightforward. When one doesn't have access to this signal, which is often the case in universal code applications of this technology, then different methods must be employed to accomplish this registration step. The example of pre-marked
10 photographic film and paper, where by definition there will never be an "unsigned" original, is a perfect case in point of the latter.

Many earlier sections have variously discussed this issue and presented certain solutions. Notably, the section on "simple" universal codes discusses one embodiment of a solution where a given master code signal is known a priori, but its precise location (and indeed, its existence or
15 non-existence) is not known. That particular section went on to give a specific example of how a very low level designed signal can be embedded within a much larger signal, wherein this designed signal is standardized so that detection equipment or reading processes can search for this standardized signal even though its exact location is unknown. The brief section on 2D universal codes went on to point out that this basic concept could be extended into 2 dimensions,
20 or, effectively, into imagery and motion pictures. Also, the section on symmetric patterns and noise patterns outlined yet another approach to the two dimensional case, wherein the nuances associated with two dimensional scale and rotation were more explicitly addressed. Therein, the idea was not merely to determine the proper orientation and scale of underlying noise patterns, but to have information transmitted as well, e.g., the N rings for the N-bit identification word.

25 This section now attempts to isolate the sub-problem of registering embedded patterns for registration's sake. Once embedded patterns are registered, we can then look again at how this registration can serve broader needs. This section presents yet another technique for embedding patterns, a technique which can be referred to as "subliminal digital graticules". "Graticules" - other words such as fiducials or reticles or hash marks could just as well be used -conveys the

- 99 -

idea of calibration marks used for the purposes of locating and/or measuring something. In this case, they are employed as low level patterns which serve as a kind of gridding function. That gridding function itself can be a carrier of 1 bit of information, as in the universal 1 second of noise (its absence or presence, copy me, don't copy me), or it can simply find the orientation and scale of other information, such as embedded signatures, or it can simply orient an image or audio object itself.

Figs. 29 and 30 visually summarize two related methods which illustrate applicant's subliminal digital graticules. As will be discussed, the method of Fig. 29 may have slight practical advantages over the method outlined in Fig. 30, but both methods effectively decompose the problem of finding the orientation of an image into a series of steps which converge on a solution. The problem as a whole can be simply stated as the following: given an arbitrary image wherein a subliminal digital graticule may have been stamped, then find the scale, rotation, and origin (offset) of the subliminal digital graticule.

The beginning point for subliminal graticules is in defining what they are. Simply put, they are visual patterns which are directly added into other images, or as the case may be, exposed onto photographic film or paper. The classic double exposure is not a bad analogy, though in digital imaging this specific concept becomes rather stretched. These patterns will generally be at a very low brightness level or exposure level, such that when they are combined with "normal" images and exposures, they will effectively be invisible (subliminal) and just as the case with embedded signatures, they will by definition not interfere with the broad value of the images to which they are added.

Figs. 29 and 30 define two classes of subliminal graticules, each as represented in the spatial frequency domain, also known as the UV plane, 1000. Common two dimensional fourier transform algorithms can transform any given image into its UV plane conjugate. To be precise, the depictions in Figs. 29 and 30 are the magnitudes of the spatial frequencies, whereas it is difficult to depict the phase and magnitude which exists at every point.

Fig. 29 shows the example of six spots in each quadrant along the 45 degree lines, 1002. These are exaggerated in this figure, in that these spots would be difficult to discern by visual inspection of the UV plane image. A rough depiction of a "typical" power spectrum of an arbitrary image as also shown, 1004. This power spectrum is generally as unique as images are
5 unique. The subliminal graticules are essentially these spots. In this example, there are six spatial frequencies combined along each of the two 45 degree axes. The magnitudes of the six frequencies can be the same or different (we'll touch upon this refinement later). Generally speaking, the phases of each are different from the others, including the phases of one 45 degree axis relative to the other. Fig. 31 depicts this graphically. The phases in this example are simply
10 randomly placed between π and $-\pi$, 1008 and 1010. Only two axes are represented in Fig. 31 - as opposed to the four separate quadrants, since the phase of the mirrored quadrants are simply $\pi/2$ out of phase with their mirrored counterparts. If we turned up the intensity on this subliminal graticule, and we transformed the result into the image domain, then we would see a weave-like cross-hatching pattern as related in the caption of Fig. 29. As stated, this weave-like
15 pattern would be at a very low intensity when added to a given image. The exact frequencies and phases of the spectral components utilized would be stored and standardized. These will become the "spectral signatures" that registration equipment and reading processes will seek to measure.

Briefly, Fig. 30 has a variation on this same general theme. Fig. 30 lays out a different class of graticules in that the spectral signature is a simple series of concentric rings rather than
20 spots along the 45 degree axes. Fig. 32 then depicts a quasi-random phase profile as a function along a half-circle (the other half of the circle then being $\pi/2$ out of phase with the first half). These are simple examples and there are a wide variety of variations possible in designing the phase profiles and the radii of the concentric rings. The transform of this type of subliminal graticule is less of a "pattern" as with the weave-like graticule of Fig. 29, where it has more of a
25 random appearance like a snowy image.

The idea behind both types of graticules is the following: embed a unique pattern into an image which virtually always will be quite distinct from the imagery into which it will be added, but which has certain properties which facilitate fast location of the pattern, as well as accuracy properties such that when the pattern is generally located, its precise location and orientation can
30 be found to some high level of precision. A corollary to the above is to design the pattern such

- 101 -

that the pattern on average minimally interferes with the typical imagery into which it will be added, and has maximum energy relative to the visibility of the pattern.

Moving on to the gross summary of how the whole process works, the graticule type of Fig. 29 facilitates an image processing search which begins by first locating the rotation axes of the subliminal graticule, then locating the scale of the graticule, then determining the origin or offset. The last step here identifies which axes is which of the two 45 degree axes by determining phase. Thus even if the image is largely upside down, an accurate determination can be made. The first step and the second step can both be accomplished using only the power spectrum data, as opposed to the phase and magnitude. The phase and magnitude signals can then be used to "fine tune" the search for the correct rotation angle and scale. The graticule of Fig. 30 switches the first two steps above, where the scale is found first, then the rotation, followed by precise determination of the origin. Those skilled in the art will recognize that determining these outstanding parameters, along two axes, are sufficient to fully register an image. The "engineering optimization challenge" is to maximize the uniqueness and brightness of the patterns relative to their visibility, while minimizing the computational overhead in reaching some specified level of accuracy and precision in registration. In the case of exposing photographic film and paper, clearly an additional engineering challenge is the outlining of economic steps to get the patterns exposed onto the film and paper in the first place, a challenge which has been addressed in previous sections.

The problem and solution as above defined is what was meant by registration for registration's sake. It should be noted that there was no mention made of making some kind of value judgement on whether or not a graticule is indeed being found or not. Clearly, the above steps could be applied to images which do not in fact have graticules inside them, the measurements then simply chasing noise. Sympathy needs to be extended to the engineer who is assigned the task of setting "detection thresholds" for these types of patterns, or any others, amidst the incredibly broad range of imagery and environmental conditions in which the patterns must be sought and verified. [Ironically, this is where the pure universal one second of noise stood in a previous section, and that was why it was appropriate to go beyond merely detecting or not detecting this singular signal, i.e. adding additional information planes]. Herein is where some real beauty shows up: in the combination of the subliminal graticules with the now-registered

- 102 -

embedded signatures described in other parts of this disclosure. Specifically, once a "candidate registration" is found - paying due homage to the idea that one may be chasing noise - then the next logical step is to perform a reading process for, e.g., a 64 bit universal code signature. As further example, we can imagine that 44 bits of the 64 bit identification word are assigned as an index of registered users -- serial numbers if you will. The remaining 20 bits are reserved as a hash code - as is well known in encryption arts - on the 44 bit identification code thus found. Thus, in one swoop, the 20 bits serve as the "yes, I have a registered image" or "no, I don't" answer. More importantly, perhaps, this allows for a system which can allow for maximum flexibility in precisely defining the levels of "false positives" in any given automated identification system. Threshold based detection will always be at the mercy of a plethora of conditions and situations, ultimately resting on arbitrary decisions. Give me N coin flips any day.

Back on point, these graticule patterns must first be added to an image, or exposed onto a piece of film. An exemplary program reads in an arbitrarily sized digital image and adds a specified graticule to the digital image to produce an output image. In the case of film, the graticule pattern would be physically exposed onto the film either before, during, or after exposure of the primary image. All of these methods have wide variations in how they might be accomplished.

The searching and registering of subliminal graticules is the more interesting and involved process. This section will first describe the elements of this process, culminating in the generalized flow chart of Fig. 37.

Fig. 33 depicts the first major "search" step in the registration process for graticules of the type in Fig. 29. A suspect image (or a scan of a suspect photograph) is first transformed in its fourier representation using well known 2D FFT routines. The input image may look like the one in Fig. 36, upper left image. Fig. 33 conceptually represents the case where the image and hence the graticules have not been rotated, though the following process fully copes with rotation issues. After the suspect image has been transformed, the power spectrum of the transform is then calculated, being simply the square root of the addition of the two squared moduli. It is also a good idea to perform a mild low pass filter operation, such as a 3x3 blur filter, on the resulting power spectrum data, so that later search steps don't need incredibly fine spaced steps. Then the

- 103 -

candidate rotation angles from 0 through 90 degrees (or 0 to $\pi/2$ in radian) are stepped through. Along any given angle, two resultant vectors are calculated, the first is the simple addition of power spectrum values at a given radius along the four lines emanating from the origin in each quadrant. The second vector is the moving average of the first vector. Then, a normalized power profile is calculated as depicted in both 1022 and 1024, the difference being that one plot is along an angle which does not align with the subliminal graticules, and the other plot does align. The normalization stipulates that the first vector is the numerator and the second vector is the denominator in the resultant vector. As can be seen in Fig. 33, 1022 and 1024, a series of peaks (which should be "six" instead of "five" as is drawn) develops when the angle aligns along its proper direction. Detection of these peaks can be effected by setting some threshold on the normalized values, and integrating their total along the whole radial line. A plot, 1026, from 0 to 90 degrees is depicted in the bottom of Fig. 33, showing that the angle 45 degrees contains the most energy. In practice, this signal is often much lower than that depicted in this bottom figure, and instead of picking the highest value as the "found rotation angle," one can simply find the top few candidate angles and submit these candidates to the next stages in the process of determining the registration. It can be appreciated by those practiced in the art that the foregoing was simply a known signal detection scheme, and that there are dozens of such schemes that can ultimately be created or borrowed. The simple requirement of the first stage process is to whittle down the candidate rotation angles to just a few, wherein more refined searches can then take over.

Fig. 34 essentially outlines the same type of gross searching in the power spectral domain. Here instead we first search for the gross scale of the concentric rings, stepping from a small scale through a large scale, rather than the rotation angle. The graph depicted in 1032 is the same normalized vector as in 1022 and 1024, but now the vector values are plotted as a function of angle around a semi-circle. The moving average denominator still needs to be calculated in the radial direction, rather than the tangential direction. As can be seen, a similar "peaking" in the normalized signal occurs when the scanned circle coincides with a graticule circle, giving rise to the plot 1040. The scale can then be found on the bottom plot by matching the known characteristics of the concentric rings (i.e. their radii) with the profile in 1040.

Fig. 35 depicts the second primary step in registering subliminal graticules of the type in Fig. 29. Once we have found a few rotation candidates from the methods of Fig. 33, we then

- 104 -

take the plots of the candidate angles of the type of 1022 and 1024 and perform what the inventor refers to as a "scaled kernel" matched filtering operation on those vectors. The scaled kernel refers to the fact that the kernel in this case is a known non-harmonic relationship of frequencies, represented as the lines with x's at the top in 1042 and 1044, and that the scale of these

5 frequencies is swept through some pre-determined range, such as 25% to 400% of some expected scale at 100%. The matched filter operation simply adds the resultant multiplied values of the scaled frequencies and their plot counterparts. Those practiced in the art will recognize the similarity of this operation with the very well known matched filter operation. The resulting plot of the matched filter operation will look something like 1046 at the bottom of Fig. 35. Each

10 candidate angle from the first step will generate its own such plot, and at this point the highest value of all of the plots will become our candidate scale, and the angle corresponding to the highest value will become our primary candidate rotation angle. Likewise for graticules of the type in Fig. 30, a similar "scaled-kernel" matched filtering operation is performed on the plot 1040 of Fig. 34. This generally provides for a single candidate scale factor. Then, using the

15 stored phase plots 1012, 1014 and 1016 of Fig. 32, a more traditional matched filtering operation is applied between these stored plots (as kernels), and the measured phase profiles along the half-circles at the previously found scale.

The last step in registering graticules of the type in Fig. 29 is to perform a garden variety matched filter between the known (either spectral or spatial) profile of the graticule with the

20 suspect image. Since both the rotation, scale and orientation are now known from previous steps, this matched filtering operation is straightforward. If the accuracies and precision of preceding steps have not exceeded design specifications in the process, then a simple micro-search can be performed in the small neighborhood about the two parameters scale and rotation, a matched filter operation performed, and the highest value found will determine a "fine tuned" scale and rotation.

25 In this way, the scale and rotation can be found to within the degree set by the noise and cross-talk of the suspect image itself. Likewise, once the scale and rotation of the graticules of the type in Fig. 30 are found, then a straightforward matched filter operation can complete the registration process, and similar "fine tuning" can be applied.

Moving on to a variant of the use of the graticules of the type in Fig. 29, Fig. 36 presents

30 the possibility for finding the subliminal graticules without the need for performing a

- 105 -

computationally expensive 2 dimensional FFT (fast fourier transform). In situations where computational overhead is a major issue, then the search problem can be reduced to a series of one-dimensional steps. Fig. 36 broadly depicts how to do this. The figure at the top left is an arbitrary image in which the graticules of the type of Fig. 29 have been embedded. Starting at angle 0, and finishing with an angle just below 180 degrees, and stepping by, for example 5 degrees, the grey values along the depicted columns can be simply added to create a resulting column-integrated scan, 1058. The figure in the top right, 1052, depicts one of the many angles at which this will be performed. This column-integrated scan then is transformed into its fourier representation using the less computationally expensive 1 dimensional FFT. This is then turned into a magnitude or "power" plot (the two are different), and a similar normalized vector version created just like 1022 and 1024 in Fig. 33. The difference now is that as the angle approaches the proper angles of the graticules, slowly the tell-tale peaks begin to appear in the 1024-like plots, but they generally show up at higher frequencies than expected for a given scale, since we are generally slightly off on our rotation. It remains to find the angle which maximizes the peak signals, thus zooming in on the proper rotation. Once the proper rotation is found, then the scaled kernel matched filter process can be applied, followed by traditional matched filtering, all as previously described. Again, the sole idea of the "short-cut" of Fig. 36 is to greatly reduce the computational overhead in using the graticules of the type in Fig. 29. The inventor has not reduced this method of Fig. 36 to practice and currently has no data on precisely how much computational savings, if any, will be realized. These efforts will be part of application specific development of the method.

Fig. 37 simply summarizes, in order of major process steps, the methods revolving around the graticules of the type in Fig. 29.

In another variant embodiment, the graticule energy is not concentrated around the 45 degree angles in the spatial frequency domain. (Some compression algorithms, such as JPEG, tend to particularly attenuate image energy at this orientation.) Instead, the energy is more widely spatially spread. Fig. 29A shows one such distribution. The frequencies near the axes, and near the origin are generally avoided, since this is where the image energy is most likely concentrated.

- 106 -

Detection of this energy in a suspect image again relies on techniques like that reviewed above. However, instead of first identifying the axes, then the rotation, and then the scale, a more global pattern matching procedure is performed in which all are determined in a brute force effort. Those skilled in the art will recognize that the Fourier-Mellin transform is well suited for
5 use in such pattern matching problems.

The foregoing principles find application, for example, in photo-duplication kiosks. Such devices typically include a lens for imaging a customer-provided original (e.g. a photographic print or film) onto an opto-electronic detector, and a print-writing device for exposing and developing an emulsion substrate (again photographic paper or film) in accordance with the image
10 data gathered by the detector. The details of such devices are well known to those skilled in the art, and are not belabored here.

In such systems, a memory stores data from the detector, and a processor (e.g. a Pentium microprocessor with associated support components) can be used to process the memory data to detect the presence of copyright data steganographically encoded therein. If such data is detected,
15 the print-writing is interrupted.

To avoid defeat of the system by manual rotation of the original image off-axis, the processor desirably implements the above-described technique to effect automatic registration of the original, notwithstanding scale, rotation, and origin offset factors. If desired, a digital signal processing board can be employed to offload certain of the FFT processing from the main (e.g.
20 Pentium) processor. After a rotated/scaled image is registered, detection of any steganographically encoded copyright notice is straightforward and assures the machine will not be used in violation of a photographer's copyright.

While the techniques disclosed above have make use of applicant's preferred steganographic encoding methods, the principles thereof are more widely applicable and can be
25 used in many instances in which automatic registration of an image is to be effected.

USE OF EMBEDDED SIGNATURES IN VIDEO, WHEREIN A VIDEO DATA STREAM
EFFECTIVELY SERVES AS A HIGH-SPEED ONE WAY MODEM

Through use of the universal coding system outlined in earlier sections, and through use of master snowy frames which change in a simple fashion frame to frame, a simple receiver can be designed such that it has pre-knowledge of the changes in the master snowy frames and can therefore read a changing N-bit message word frame by frame (or I-frame by I-frame as the case may be in MPEG video). In this way, a motion picture sequence can be used as a high speed one-way information channel, much like a one-way modem. Consider, for example, a frame of video data with N scan lines which is steganographically encoded to effect the transmission of an N-bit message. If there are 484 scan lines in a frame (N), and frames change 30 times a second, an information channel with a capacity comparable to a 14.4 kilobaud modem is achieved.

In actual practice, a data rate substantially in excess of N bits per frame can usually be achieved, yielding transmission rates nearer that of ISDN circuits.

15 FRAUD PREVENTION IN WIRELESS COMMUNICATIONS

In the cellular telephone industry, hundreds of millions of dollars of revenue is lost each year through theft of services. While some services are lost due to physical theft of cellular telephones, the more pernicious threat is posed by cellular telephone hackers.

Cellular telephone hackers employ various electronic devices to mimic the identification signals produced by an authorized cellular telephone. (These signals are sometimes called authorization signals, verification numbers, signature data, etc.) Often, the hacker learns of these signals by eavesdropping on authorized cellular telephone subscribers and recording the data exchanged with the cell cite. By artful use of this data, the hacker can impersonate an authorized subscriber and dupe the carrier into completing pirate calls.

- 108 -

In the prior art, identification signals are segregated from the voice signals. Most commonly, they are temporally separated, e.g. transmitted in a burst at the time of call origination. Voice data passes through the channel only after a verification operation has taken place on this identification data. (Identification data is also commonly included in data packets sent during the transmission.) Another approach is to spectrally separate the identification, e.g. in a spectral subband outside that allocated to the voice data.

Other fraud-deterrent schemes have also been employed. One class of techniques monitors characteristics of a cellular telephone's RF signal to identify the originating phone. Another class of techniques uses handshaking protocols, wherein some of the data returned by the cellular telephone is based on an algorithm (e.g. hashing) applied to random data sent thereto.

Combinations of the foregoing approaches are also sometimes employed.

U.S. Patents 5,465,387, 5,454,027, 5,420,910, 5,448,760, 5,335,278, 5,345,595, 5,144,649, 5,204,902, 5,153,919 and 5,388,212 detail various cellular telephone systems, and fraud deterrence techniques used therein. The disclosures of these patents are incorporated by reference.

As the sophistication of fraud deterrence systems increases, so does the sophistication of cellular telephone hackers. Ultimately, hackers have the upper hand since they recognize that all prior art systems are vulnerable to the same weakness: the identification is based on some attribute of the cellular telephone transmission outside the voice data. Since this attribute is segregated from the voice data, such systems will always be susceptible to pirates who electronically "patch" their voice into a composite electronic signal having the attribute(s) necessary to defeat the fraud deterrence system.

To overcome this failing, preferred embodiments of this aspect of the present technology steganographically encode the voice signal with identification data, resulting in "in-band" signalling (in-band both temporally and spectrally). This approach allows the carrier to monitor the user's voice signal and decode the identification data therefrom.

- 109 -

In one such form of the technology, some or all of the identification data used in the prior art (e.g. data transmitted at call origination) is repeatedly steganographically encoded in the user's voice signal as well. The carrier can thus periodically or aperiodically check the identification data accompanying the voice data with that sent at call origination to ensure they match. If they do not, the call is identified as being hacked and steps for remediation can be instigated such as interrupting the call.

In another form of the technology, a randomly selected one of several possible messages is repeatedly steganographically encoded on the subscriber's voice. An index sent to the cellular carrier at call set-up identifies which message to expect. If the message steganographically decoded by the cellular carrier from the subscriber's voice does not match that expected, the call is identified as fraudulent.

In a preferred form of this aspect of the technology, the steganographic encoding relies on a pseudo random data signal to transform the message or identification data into a low level noise-like signal superimposed on the subscriber's digitized voice signal. This pseudo random data signal is known, or knowable, to both the subscriber's telephone (for encoding) and to the cellular carrier (for decoding). Many such embodiments rely on a deterministic pseudo random number generator seeded with a datum known to both the telephone and the carrier. In simple embodiments this seed can remain constant from one call to the next (e.g. a telephone ID number). In more complex embodiments, a pseudo-one-time pad system may be used, wherein a new seed is used for each session (i.e. telephone call). In a hybrid system, the telephone and cellular carrier each have a reference noise key (e.g. 10,000 bits) from which the telephone selects a field of bits, such as 50 bits beginning at a randomly selected offset, and each uses this excerpt as the seed to generate the pseudo random data for encoding. Data sent from the telephone to the carrier (e.g. the offset) during call set-up allows the carrier to reconstruct the same pseudo random data for use in decoding. Yet further improvements can be derived by borrowing basic techniques from the art of cryptographic communications and applying them to the steganographically encoded signal detailed in this disclosure.

- 110 -

Details of applicant's preferred techniques for steganographic encoding/decoding with a pseudo random data stream are more particularly detailed in the previous portions of this specification, but the present technology is not limited to use with such techniques.

The reader is presumed to be familiar with cellular communications technologies.

5 Accordingly, details known from prior art in this field aren't belabored herein.

Referring to Fig. 38, an illustrative cellular system includes a telephone 2010, a cell site 2012, and a central office 2014.

Conceptually, the telephone may be viewed as including a microphone 2016, an A/D converter 2018, a data formatter 2020, a modulator 2022, an RF section 2024, an antenna 2026, a
10 demodulator 2028, a data unformatter 2030, a D/A converter 2032, and a speaker 2034.

In operation, a subscriber's voice is picked up by the microphone 2016 and converted to digital form by the A/D converter 2018. The data formatter 2020 puts the digitized voice into packet form, adding synchronization and control bits thereto. The modulator 2022 converts this digital data stream into an analog signal whose phase and/or amplitude properties change in
15 accordance with the data being modulated. The RF section 2024 commonly translates this time-varying signal to one or more intermediate frequencies, and finally to a UHF transmission frequency. The RF section thereafter amplifies it and provides the resulting signal to the antenna 2026 for broadcast to the cell site 2012.

The process works in reverse when receiving. A broadcast from the cell site is received
20 through the antenna 2026. RF section 2024 amplifies and translates the received signal to a different frequency for demodulation. Demodulator 2028 processes the amplitude and/or phase variations of the signal provided by the RF section to produce a digital data stream corresponding thereto. The data unformatter 2030 segregates the voice data from the associated synchronization/control data, and passes the voice data to the D/A converter for conversion into
25 analog form. The output from the D/A converter drives the speaker 2034, through which the subscriber hears the other party's voice.

- 111 -

The cell site 2012 receives broadcasts from a plurality of telephones 2010, and relays the data received to the central office 2014. Likewise, the cell site 2012 receives outgoing data from the central office and broadcasts same to the telephones.

The central office 2014 performs a variety of operations, including call authentication,
5 switching, and cell hand-off.

(In some systems, the functional division between the cell site and the central station is different than that outlined above. Indeed, in some systems, all of this functionality is provided at a single site.)

In an exemplary embodiment of this aspect of the technology, each telephone 2010
10 additionally includes a steganographic encoder 2036. Likewise, each cell site 2012 includes a steganographic decoder 2038. The encoder operates to hide an auxiliary data signal among the signals representing the subscriber's voice. The decoder performs the reciprocal function, discerning the auxiliary data signal from the encoded voice signal. The auxiliary signal serves to verify the legitimacy of the call.

15 An exemplary steganographic encoder 2036 is shown in Fig. 39.

The illustrated encoder 2036 operates on digitized voice data, auxiliary data, and pseudo-random noise (PRN) data. The digitized voice data is applied at a port 2040 and is provided, e.g., from A/D converter 2018. The digitized voice may comprise 8-bit samples. The auxiliary data is applied at a port 2042 and comprises, in one form of the technology, a stream of binary
20 data uniquely identifying the telephone 2010. (The auxiliary data may additionally include administrative data of the sort conventionally exchanged with a cell site at call set-up.) The pseudo-random noise data is applied at a port 2044 and can be, e.g., a signal that randomly alternates between "-1" and "1" values. (More and more cellular phones are incorporating spread spectrum capable circuitry, and this pseudo-random noise signal and other aspects of this
25 technology can often "piggy-back" or share the circuitry which is already being applied in the basic operation of a cellular unit).

- 112 -

For expository convenience, it is assumed that all three data signals applied to the encoder 2036 are clocked at a common rate, although this is not necessary in practice.

In operation, the auxiliary data and PRN data streams are applied to the two inputs of a logic circuit 2046. The output of circuit 2046 switches between -1 and +1 in accordance with the following table:

| AUX | PRN | OUTPUT |
|-----|-----|--------|
| 0 | -1 | 1 |
| 0 | 1 | -1 |
| 1 | -1 | -1 |
| 1 | 1 | 1 |

(If the auxiliary data signal is conceptualized as switching between -1 and 1, instead of 0 and 1, it will be seen that circuit 2046 operates as a one-bit multiplier.)

The output from gate 2046 is thus a bipolar data stream whose instantaneous value changes randomly in accordance with the corresponding values of the auxiliary data and the PRN data. It may be regarded as noise. However, it has the auxiliary data encoded therein. The auxiliary data can be extracted if the corresponding PRN data is known.

The noise-like signal from gate 2046 is applied to the input of a scaler circuit 2048. Scaler circuit scales (e.g. multiplies) this input signal by a factor set by a gain control circuit 2050. In the illustrated embodiment, this factor can range between 0 and 15. The output from scaler circuit 2048 can thus be represented as a five-bit data word (four bits, plus a sign bit) which changes each clock cycle, in accordance with the auxiliary and PRN data, and the scale

- 113 -

factor. The output from the scaler circuit may be regarded as "scaled noise data" (but again it is "noise" from which the auxiliary data can be recovered, given the PRN data).

The scaled noise data is summed with the digitized voice data by a summer 2051 to provide the encoded output signal (e.g. binarily added on a sample by sample basis). This output
5 signal is a composite signal representing both the digitized voice data and the auxiliary data.

The gain control circuit 2050 controls the magnitude of the added scaled noise data so its addition to the digitized voice data does not noticeably degrade the voice data when converted to analog form and heard by a subscriber. The gain control circuit can operate in a variety of ways.

One is a logarithmic scaling function. Thus, for example, voice data samples having
10 decimal values of 0, 1 or 2 may be correspond to scale factors of unity, or even zero, whereas voice data samples having values in excess of 200 may correspond to scale factors of 15. Generally speaking, the scale factors and the voice data values correspond by a square root relation. That is, a four-fold increase in a value of the voice data corresponds to approximately a two-fold increase in a value of the scaling factor associated therewith. Another scaling function
15 would be linear as derived from the average power of the voice signal.

(The parenthetical reference to zero as a scaling factor alludes to cases, e.g., in which the digitized voice signal sample is essentially devoid of information content.)

More satisfactory than basing the instantaneous scaling factor on a single voice data sample, is to base the scaling factor on the dynamics of several samples. That is, a stream of
20 digitized voice data which is changing rapidly can camouflage relatively more auxiliary data than a stream of digitized voice data which is changing slowly. Accordingly, the gain control circuit 2050 can be made responsive to the first, or preferably the second- or higher-order derivative of the voice data in setting the scaling factor.

In still other embodiments, the gain control block 2050 and scaler 2048 can be omitted
25 entirely.

- 114 -

(Those skilled in the art will recognize the potential for "rail errors" in the foregoing systems. For example, if the digitized voice data consists of 8-bit samples, and the samples span the entire range from 0 to 255 (decimal), then the addition or subtraction of scaled noise to/from the input signal may produce output signals that cannot be represented by 8 bits (e.g. -2, or 257).

5 A number of well-understood techniques exist to rectify this situation, some of them proactive and some of them reactive. Among these known techniques are: specifying that the digitized voice data shall not have samples in the range of 0-4 or 241-255, thereby safely permitting combination with the scaled noise signal; and including provision for detecting and adaptively modifying digitized voice samples that would otherwise cause rail errors.)

10 Returning to the telephone 2010, an encoder 2036 like that detailed above is desirably interposed between the A/D converter 2018 and the data formatter 2020, thereby serving to steganographically encode all voice transmissions with the auxiliary data. Moreover, the circuitry or software controlling operation of the telephone is arranged so that the auxiliary data is encoded repeatedly. That is, when all bits of the auxiliary data have been encoded, a pointer loops back
15 and causes the auxiliary data to be applied to the encoder 2036 anew. (The auxiliary data may be stored at a known address in RAM memory for ease of reference.)

It will be recognized that the auxiliary data in the illustrated embodiment is transmitted at a rate one-eighth that of the voice data. That is, for every 8-bit sample of voice data, scaled noise data corresponding to a single bit of the auxiliary data is sent. Thus, if voice samples are
20 sent at a rate of 4800 samples/second, auxiliary data can be sent at a rate of 4800 bits/second. If the auxiliary data is comprised of 8-bit symbols, auxiliary data can be conveyed at a rate of 600 symbols/second. If the auxiliary data consists of a string of even 60 symbols, each second of voice conveys the auxiliary data ten times. (Significantly higher auxiliary data rates can be achieved by resorting to more efficient coding techniques, such as limited-symbol codes (e.g. 5-
25 or 6-bit codes), Huffman coding, etc.) This highly redundant transmission of the auxiliary data permits lower amplitude scaled noise data to be used while still providing sufficient signal-to-noise headroom to assure reliable decoding -- even in the relatively noisy environment associated with radio transmissions.

- 115--

Turning now to Fig. 40, each cell site 2012 has a steganographic decoder 2038 by which it can analyze the composite data signal broadcast by the telephone 2010 to discern and separate the auxiliary data and digitized voice data therefrom. (The decoder desirably works on unformatted data (i.e. data with the packet overhead, control and administrative bits removed; this is not shown for clarity of illustration).

The decoding of an unknown embedded signal (i.e. the encoded auxiliary signal) from an unknown voice signal is best done by some form of statistical analysis of the composite data signal. The techniques therefor discussed above are equally applicable here. For example, the entropy-based approach can be utilized. In this case, the auxiliary data repeats every 480 bits (instead of every 8). As above, the entropy-based decoding process treats every 480th sample of the composite signal in like fashion. In particular, the process begins by adding to the 1st, 481st, 861st, etc. samples of the composite data signal the PRN data with which these samples were encoded. (That is, a set of sparse PRN data is added: the original PRN set, with all but every 480th datum zeroed out.) The localized entropy of the resulting signal around these points (i.e. the composite data signal with every 480th sample modified) is then computed.

The foregoing step is then repeated, this time subtracting the PRN data corresponding thereto from the 1st, 481st, 961st, etc. composite data samples.

One of these two operations will counteract (e.g. undo) the encoding process and reduce the resulting signal's entropy; the other will aggravate it. If adding the sparse PRN data to the composite data reduces its entropy, then this data must earlier have been subtracted from the original voice signal. This indicates that the corresponding bit of the auxiliary data signal was a "0" when these samples were encoded. (A "0" at the auxiliary data input of logic circuit 46 caused it to produce an inverted version of the corresponding PRN datum as its output datum, resulting in subtraction of the corresponding PRN datum from the voice signal.)

Conversely, if subtracting the sparse PRN data from the composite data reduces its entropy, then the encoding process must have earlier added this noise. This indicates that the value of the auxiliary data bit was a "1" when samples 1, 481, 961, etc., were encoded.

- 116 -

By noting in which case entropy is lower by (a) adding or (b) subtracting a sparse set of PRN data to/from the composite data, it can be determined whether the first bit of the auxiliary data is (a) a "0", or (b) a "1." (In real life applications, in the presence of various distorting phenomena, the composite signal may be sufficiently corrupted so that neither adding nor subtracting the sparse PRN data actually reduces entropy. Instead, both operations will increase entropy. In this case, the "correct" operation can be discerned by observing which operation increases the entropy less.)

The foregoing operations can then be conducted for the group of spaced samples of the composite data beginning with the second sample (i.e. 2, 482, 962, ...). The entropy of the resulting signals indicate whether the second bit of the auxiliary data signal is a "0" or a "1." Likewise with the following 478 groups of spaced samples in the composite signal, until all 480 bits of the code word have been discerned.

As discussed above, correlation between the composite data signal and the PRN data can also be used as a statistical detection technique. Such operations are facilitated in the present context since the auxiliary data whose encoded representation is sought, is known, at least in large part, a priori. (In one form of the technology, the auxiliary data is based on the authentication data exchanged at call set-up, which the cellular system has already received and logged; in another form (detailed below), the auxiliary data comprises a predetermined message.) Thus, the problem can be reduced to determining whether an expected signal is present or not (rather than looking for an entirely unknown signal). Moreover, data formatter 2020 breaks the composite data into frames of known length. (In a known GSM implementation, voice data is sent in time slots which convey 114 data bits each.) By padding the auxiliary data as necessary, each repetition of the auxiliary data can be made to start, e.g., at the beginning of such a frame of data. This, too, simplifies the correlation determinations, since 113 of every 114 possible bit alignments can be ignored (facilitating decoding even if none of the auxiliary data is known a priori).

Again, this wireless fraud detection poses the now-familiar problem of detecting known signals in noise, and the approaches discussed earlier are equally applicable here.

- 117 -

Where, as here, the location of the auxiliary signal is known a priori (or more accurately, known to fall within one of a few discrete locations, as discussed above), then the matched filter approach can often be reduced to a simple vector dot product between a set of sparse PRN data, and mean-removed excerpts of the composite signal corresponding thereto. (Note that the PRN data need not be sparse and may arrive in contiguous bursts, such as in British patent publication 2,196,167 mentioned earlier wherein a given bit in a message has contiguous PRN values associated with it.) Such a process steps through all 480 sparse sets of PRN data and performs corresponding dot product operations. If the dot product is positive, the corresponding bit of the auxiliary data signal is a "1;" if the dot product is negative, the corresponding bit of the auxiliary data signal is a "0." If several alignments of the auxiliary data signal within the framed composite signal are possible, this procedure is repeated at each candidate alignment, and the one yielding the highest correlation is taken as true. (Once the correct alignment is determined for a single bit of the auxiliary data signal, the alignment of all the other bits can be determined therefrom. "Alignment," perhaps better known as "synchronization," can be achieved by primarily through the very same mechanisms which lock on and track the voice signal itself and allow for the basic functioning of the cellular unit).

Security Considerations

Security of the presently described embodiments depends, in large part, on security of the PRN data and/or security of the auxiliary data. In the following discussion, a few of many possible techniques for assuring the security of these data are discussed.

In a first embodiment, each telephone 2010 is provided with a long noise key unique to the telephone. This key may be, e.g., a highly unique 10,000 bit string stored in ROM. (In most applications, keys substantially shorter than this may be used.)

The central office 2014 has access to a secure disk 2052 on which such key data for all authorized telephones are stored. (The disk may be remote from the office itself.)

- 118 -

Each time the telephone is used, fifty bits from this noise key are identified and used as the seed for a deterministic pseudo random number generator. The data generated by this PRN generator serve as the PRN data for that telephone call.

5 The fifty bit seed can be determined, e.g., by using a random number generator in the telephone to generate an offset address between 0 and 9,950 each time the telephone is used to place a call. The fifty bits in the noise key beginning at this offset address are used as the seed.

During call setup, this offset address is transmitted by the telephone, through the cell site 2012, to the central office 2014. There, a computer at the central office uses the offset address to index its copy of the noise key for that telephone. The central office thereby identifies the same
10 50 bit seed as was identified at the telephone. The central office 2014 then relays these 50 bits to the cell site 2012, where a deterministic noise generator like that in the telephone generates a PRN sequence corresponding to the 50 bit key and applies same to its decoder 2038.

By the foregoing process, the same sequence of PRN data is generated both at the telephone and at the cell site. Accordingly, the auxiliary data encoded on the voice data by the
15 telephone can be securely transmitted to, and accurately decoded by, the cell site. If this auxiliary data does not match the expected auxiliary data (e.g. data transmitted at call set-up), the call is flagged as fraudulent and appropriate remedial action is taken.

It will be recognized that an eavesdropper listening to radio transmission of call set-up information can intercept only the randomly generated offset address transmitted by the telephone
20 to the cell site. This data, alone, is useless in pirating calls. Even if the hacker had access to the signals provided from the central office to the cell site, this data too is essentially useless: all that is provided is a 50 bit seed. Since this seed is different for nearly each call (repeating only 1 out of every 9,950 calls), it too is unavailing to the hacker.

In a related system, the entire 10,000 bit noise key can be used as a seed. An offset
25 address randomly generated by the telephone during call set-up can be used to identify where, in the PRN data resulting from that seed, the PRN data to be used for that session is to begin.

- 119 -

(Assuming 4800 voice samples per second, 4800 PRN data are required per second, or about 17 million PRN data per hour. Accordingly, the offset address in this variant embodiment will likely be far larger than the offset address described above.)

5 In this variant embodiment, the PRN data used for decoding is preferably generated at the central station from the 10,000 bit seed, and relayed to the cell site. (For security reasons, the 10,000 bit noise key should not leave the security of the central office.)

In variants of the foregoing systems, the offset address can be generated by the central station or at the cell site, and relayed to the telephone during call set-up, rather than vice versa.

10 In another embodiment, the telephone 2010 may be provided with a list of one-time seeds, matching a list of seeds stored on the secure disk 2052 at the central office. Each time the telephone is used to originate a new call, the next seed in the list is used. By this arrangement, no data needs to be exchanged relating to the seed; the telephone and the carrier each independently know which seed to use to generate the pseudo random data sequence for the current session.

15 In such an embodiment, the carrier can determine when the telephone has nearly exhausted its list of seeds, and can transmit a substitute list (e.g. as part of administrative data occasionally provided to the telephone). To enhance security, the carrier may require that the telephone be returned for manual reprogramming, to avoid radio transmission of this sensitive information. Alternatively, the substitute seed list can be encrypted for radio transmission using
20 any of a variety of well known techniques.

In a second class of embodiments, security derives not from the security of the PRN data, but from security of the auxiliary message data encoded thereby. One such system relies on transmission of a randomly selected one of 256 possible messages.

25 In this embodiment, a ROM in the telephone stores 256 different messages (each message may be, e.g., 128 bits in length). When the telephone is operated to initiate a call, the telephone

- 120 -

randomly generates a number between 1 and 256, which serves as an index to these stored messages. This index is transmitted to the cell site during call set-up, allowing the central station to identify the expected message from a matching database on secure disk 2052 containing the same 256 messages. (Each telephone has a different collection of messages.) (Alternatively, the carrier may randomly select the index number during call set-up and transmit it to the telephone, identifying the message to be used during that session.) In a theoretically pure world where proposed attacks to a secure system are only mathematical in nature, much of these additional layers of security might seem superfluous. (The addition of these extra layers of security, such as differing the messages themselves, simply acknowledge that the designer of actual public-functioning secure systems will face certain implementation economics which might compromise the mathematical security of the core principles of this technology, and thus these auxiliary layers of security may afford new tools against the inevitable attacks on implementation).

Thereafter, all voice data transmitted by the telephone for the duration of that call is steganographically encoded with the indexed message. The cell site checks the data received from the telephone for the presence of the expected message. If the message is absent, or if a different message is decoded instead, the call is flagged as fraudulent and remedial action is taken.

In this second embodiment, the PRN data used for encoding and decoding can be as simple or complex as desired. A simple system may use the same PRN data for each call. Such data may be generated, e.g., by a deterministic PRN generator seeded with fixed data unique to the telephone and known also by the central station (e.g. a telephone identifier), or a universal noise sequence can be used (i.e. the same noise sequence can be used for all telephones). Or the pseudo random data can be generated by a deterministic PRN generator seeded with data that changes from call to call (e.g. based on data transmitted during call set-up identifying, e.g., the destination telephone number, etc.). Some embodiments may seed the pseudo random number generator with data from a preceding call (since this data is necessarily known to the telephone and the carrier, but is likely not known to pirates).

Naturally, elements from the foregoing two approaches can be combined in various ways, and supplemented by other features. The foregoing embodiments are exemplary only, and do not

- 121 -

begin to catalog the myriad approaches which may be used. Generally speaking, any data which is necessarily known or knowable by both the telephone and the cell site/central station, can be used as the basis for either the auxiliary message data, or the PRN data by which it is encoded.

Since the preferred embodiments of this aspect of the present technology each redundantly
5 encodes the auxiliary data throughout the duration of the subscriber's digitized voice, the auxiliary data can be decoded from any brief sample of received audio. In the preferred forms of this aspect of the technology, the carrier repeatedly checks the steganographically encoded auxiliary data (e.g. every 10 seconds, or at random intervals) to assure that it continues to have the expected attributes.

10 While the foregoing discussion has focused on steganographically encoding a transmission from a cellular telephone, it will be recognized that transmissions to a cellular telephone can be steganographically encoded as well. Such arrangements find applicability, e.g., in conveying administrative data (i.e. non-voice data) from the carrier to individual telephones. This administrative data can be used, for example, to reprogram parameters of targeted cellular
15 telephones (or all cellular telephones) from a central location, to update seed lists (for systems employing the above-described on-time pad system), to apprise "roaming" cellular telephones of data unique to an unfamiliar local area, etc.

In some embodiments, the carrier may steganographically transmit to the cellular telephone a seed which the cellular phone is to use in its transmissions to the carrier during the
20 remainder of that session.

While the foregoing discussion has focused on steganographic encoding of the baseband digitized voice data, artisans will recognize that intermediate frequency signals (whether analog or digital) can likewise be steganographically encoded in accordance with principles of the technology. An advantage of post-baseband encoding is that the bandwidth of these intermediate
25 signals is relatively large compared with the baseband signal, allowing more auxiliary data to be encoded therein, or allowing a fixed amount of auxiliary data to be repeated more frequently during transmission. (If steganographic encoding of an intermediate signal is employed, care

- 122 -

should be taken that the perturbations introduced by the encoding are not so large as to interfere with reliable transmission of the administrative data, taking into account any error correcting facilities supported by the packet format).

Those skilled in the art will recognize that the auxiliary data, itself, can be arranged in known ways to support error detecting, or error correcting capabilities by the decoder 38. The interested reader is referred, e.g., to Rorabaugh, *Error Coding Cookbook*, McGraw Hill, 1996, one of many readily available texts detailing such techniques.

While the preferred embodiment of this aspect of the technology is illustrated in the context of a cellular system utilizing packetized data, other wireless systems do not employ such conveniently framed data. In systems in which framing is not available as an aid to synchronization, synchronization marking can be achieved within the composite data signal by techniques such as that detailed in applicant's prior applications. In one class of such techniques, the auxiliary data itself has characteristics facilitating its synchronization. In another class of techniques, the auxiliary data modulates one or more embedded carrier patterns which are designed to facilitate alignment and detection.

As noted earlier, the principles of the technology are not restricted to use with the particular forms of steganographic encoding detailed above. Indeed, any steganographic encoding technique previously known, or hereafter invented, can be used in the fashion detailed above to enhance the security or functionality of cellular (or other wireless, e.g. PCS) communications systems. Likewise, these principles are not restricted to wireless telephones; any wireless transmission may be provided with an "in-band" channel of this type.

It will be recognized that systems for implementing applicant's technology can comprises dedicated hardware circuit elements, but more commonly comprise suitably programmed microprocessors with associated RAM and ROM memory (e.g. one such system in each of the telephone 2010, cell-site 2012, and central office 2014).

ENCODING BY BIT CELLS

The foregoing discussions have focused on incrementing or decrementing the values of individual pixels, or of groups of pixels (bumps), to reflect encoding of an auxiliary data signal combined with a pseudo random noise signal. The following discussion details a variant
5 embodiment wherein the auxiliary data -- without pseudo randomization -- is encoded by patterned groups of pixels, here termed bit cells.

Referring to Figs 41A and 41B, two illustrative 2x2 bit cells are shown. Fig. 41A is used to represent a "0" bit of the auxiliary data, while Fig. 41B is used to represent a "1" bit. In operation, the pixels of the underlying image are tweaked up or down in accordance with the +/-
10 values of the bit cells to represent one of these two bit values. (The magnitude of the tweaking at any given pixel, bit cell or region of the image can be a function of many factors, as detailed below. It is the sign of the tweaking that defines the characteristic pattern.) In decoding, the relative biases of the encoded pixels are examined (using techniques described above) to identify, for each corresponding region of the encoded image, which of the two patterns is represented.

15 While the auxiliary data is not explicitly randomized in this embodiment, it will be recognized that the bit cell patterns may be viewed as a "designed" carrier signal, as discussed above.

The substitution of the previous embodiments' pseudo random noise with the present "designed" information carrier affords an advantage: the bit cell patterning manifests itself in
20 Fourier space. Thus, the bit cell patterning can act like the subliminal digital graticules discussed above to help register a suspect image to remove scale/rotation errors. By changing the size of the bit cell, and the pattern therein, the location of the energy thereby produced in the spatial transform domain can be tailored to optimize independence from typical imagery energy and facilitate detection.

- 124 -

(While the foregoing discussion contemplates that the auxiliary data is encoded directly -- without randomization by a PRN signal, in other embodiments, randomization can of course be used.)

5 MORE ON PERCEPTUALLY ADAPTIVE SIGNING

In several of the above-detailed embodiments, the magnitude of the signature energy was tailored on a region-by-region basis to reduce its visibility in an image (or its audibility in a sound, etc.). In the following discussion, applicant more particularly considers the issue of hiding signature energy in an image, the separate issues thereby posed, and solutions to each of these
10 issues.

The goal of the signing process, beyond simply functioning, is to maximize the "numeric detectability" of an embedded signature while meeting some form of fixed "visibility/acceptability threshold" set by a given user/creator.

In service to design toward this goal, imagine the following three axis parameter space,
15 where two of the axes are only half-axes (positive only), and the third is a full axis (both negative and positive). This set of axes define two of the usual eight octal spaces of euclidean 3-space. As things refine and "deservedly separable" parameters show up on the scene (such as "extended local visibility metrics"), then they can define their own (generally) half-axis and extend the following example beyond three dimensions.

20 The signing design goal becomes optimally assigning a "gain" to a local bump based on its coordinates in the above defined space, whilst keeping in mind the basic needs of doing the operations fast in real applications. To begin with, the three axes are the following. We'll call the two half axes x and y, while the full axis will be z.

The x axis represents the luminance of the singular bump. The basic idea is that you can squeeze a little more energy into bright regions as opposed to dim ones. It is important to note that when true "psycho-linear - device independent" luminance values (pixel DN's) come along, this axis might become superfluous, unless of course if the luminance value couples into the other operative axes (e.g. $C*xy$). For now, this is here as much due to the sub-optimality of current quasi-linear luminance coding.

The y axis is the "local hiding potential" of the neighborhood within which the bump finds itself. The basic idea is that flat regions have a low hiding potential since the eye can detect subtle changes in such regions, whereas complex textured regions have a high hiding potential. Long lines and long edges tend toward the lower hiding potential since "breaks and choppiness" in nice smooth long lines are also somewhat visible, while shorter lines and edges, and mosaics thereof, tend toward the higher hiding potential. These latter notions of long and short are directly connected to processing time issues, as well to issues of the engineering resources needed to carefully quantify such parameters. Developing the working model of the y-axis will inevitably entail one part theory to one part picky-artist-empiricism. As the parts of the hodge-podge y-axis become better known, they can splinter off into their own independent axes if it's worth it.

The z-axis is the "with or against the grain" (discussed below) axis which is the full axis - as opposed to the other two half-axes. The basic idea is that a given input bump has a pre-existing bias relative to whether one wishes to encode a '1' or a '0' at its location, which to some non-trivial extent is a function of the reading algorithms which will be employed, whose (bias) magnitude is semi-correlated to the "hiding potential" of the y-axis, and, fortunately, can be used advantageously as a variable in determining what magnitude of a tweak value is assigned to the bump in question. The concomitant basic idea is that when a bump is already your friend (i.e. its bias relative to its neighbors already tends towards the desired delta value), then don't change it much. Its natural state already provides the delta energy needed for decoding, without altering the localized image value much, if at all. Conversely, if a bump is initially your enemy (i.e. its bias relative to its neighbors tends away from the delta sought to be imposed by the encoding), then change it an exaggerated amount. This later operation tends to reduce the excursion of this point relative to its neighbors, making the point less visibly conspicuous (a

- 126 -

highly localized blurring operation), while providing additional energy detectable when decoding. These two cases are termed "with the grain" and "against the grain" herein.

The above general description of the problem should suffice for many years. Clearly adding in chrominance issues will expand the definitions a bit, leading to a bit more signature
5 bang for the visibility, and human visibility research which is applied to the problem of compression can equally be applied to this area but for diametrically opposed reasons. Here are guiding principles which can be employed in an exemplary application.

For speed's sake, local hiding potential can be calculated only based on a 3 by 3 neighborhood of pixels, the center one being signed and its eight neighbors. Beyond speed
10 issues, there is also no data or coherent theory to support anything larger as well. The design issue boils down to canning the y-axis visibility thing, how to couple the luminance into this, and a little bit on the friend/enemy asymmetry thing. A guiding principle is to simply make a flat region zero, a classic pure maxima or minima region a "1.0" or the highest value, and to have "local lines", "smooth slopes", "saddle points" and whatnot fall out somewhere in between.

15 The exemplary application uses six basic parameters: 1) luminance; 2) difference from local average; 3) the asymmetry factor (with or against the grain); 4) minimum linear factor (our crude attempt at flat v. lines v. maxima); 5) bit plane bias factor; and 6) global gain (the user's single top level gain knob).

20 The Luminance, and Difference from Local Average parameters are straight forward, and their use is addressed elsewhere in this specification.

The Asymmetry factor is a single scalar applied to the "against the grain" side of the difference axis of number 2 directly above.

The Minimum Linear factor is admittedly crude but it should be of some service even in a 3 by 3 neighborhood setting. The idea is that true 2D local minima and maxima will be highly
25 perturbed along each of the four lines travelling through the center pixel of the 3 by 3

- 127 -

neighborhood, while a visual line or edge will tend to flatten out at least one of the four linear profiles. [The four linear profiles are each 3 pixels in length, i.e., the top left pixel - center - bottom right; the top center - center - bottom center; the top right - center - bottom left; the right center - center - left center;]. Let's choose some metric of entropy as applied to three pixels in a row, perform this on all four linear profiles, then choose the minimum value for our ultimate parameter to be used as our 'y-axis.'

The Bit Plane Bias factor is an interesting creature with two faces, the pre-emptive face and the post-emptive face. In the former, you simply "read" the unsigned image and see where all the biases fall out for all the bit planes, then simply boost the "global gain" of the bit planes which are, in total, going against your desired message, and leave the others alone or even slightly lower their gain. In the post-emptive case, you churn out the whole signing process replete with the pre-emptive bit plane bias and the other 5 parameters listed here, and then you, e.g., run the signed image through heavy JPEG compression AND model the "gestalt distortion" of line screen printing and subsequent scanning of the image, and then you read the image and find out which bit planes are struggling or even in error, you appropriately beef up the bit plane bias, and you run through the process again. If you have good data driving the beefing process you should only need to perform this step once, or, you can easily Van-Cittertize the process (arcane reference to reiterate the process with some damping factor applied to the tweaks).

Finally, there is the Global Gain. The goal is to make this single variable the top level "intensity knob" (more typically a slider or other control on a graphical user interface) that the slightly curious user can adjust if they want to. The very curious user can navigate down advanced menus to get their experimental hands on the other five variables here, as well as others.

Visible Watermark

- 128 -

In certain applications it is desirable to apply a visible indicia to an image to indicate that it includes steganographically encoded data. In one embodiment, this indicia can be a lightly visible logo (sometimes termed a "watermark") applied to one corner of the image. This indicates that the image is a "smart" image, conveying data in addition to the imagery. A lightbulb is one
5 suitable logo.

If marking of images becomes widespread (e.g., by software compatible with Adobe's image processing software), a user of such software can decode the embedded data from an image and consult a public registry to identify the proprietor of the image. In some embodiments, the registry can serve as the conduit through which appropriate royalty payments are forwarded to the proprietor for the user's use of an image. (In an illustrative embodiment, the registry is a server on the Internet, accessible via the World Wide Web, coupled to a database. The database includes detailed information on catalogued images (e.g. name, address, phone number of

20

- 130 -

proprietor, and a schedule of charges for different types of uses to which the image may be put), indexed by identification codes with which the images themselves are encoded. A person who decodes an image queries the registry with the codes thereby gleaned to obtain the desired data and, if appropriate, to forward electronic payment of a copyright royalty to the image's proprietor.)

Particular Data Formats

While the foregoing steganography techniques are broadly applicable, their commercial acceptance will be aided by establishment of standards setting forth which pixels/bit cells represent what. The following discussion proposes one set of possible standards. For expository convenience, this discussion focuses on decoding of the data; encoding follows in a straightforward manner.

Referring to Fig. 42, an image 1202 includes a plurality of tiled "signature blocks" 1204. (Partial signature blocks may be present at the image edges.) Each signature block 1204 includes an 8 x 8 array of sub-blocks 1206. Each sub-block 1206 includes an 8 x 8 array of bit cells 1208. Each bit cell comprises a 2 x 2 array of "bumps" 1210. Each bump 1210, in turn, comprises a square grouping of 16 individual pixels 1212.

The individual pixels 1212 are the smallest quanta of image data. In this arrangement, however, pixel values are not, individually, the data carrying elements. Instead, this role is served by bit cells 1208 (i.e. 2 x 2 arrays of bumps 1210). In particular, the bumps comprising the bits cells are encoded to assume one of the two patterns shown in Fig. 41. As noted earlier, the pattern shown in Fig. 41A represents a "0" bit, while the pattern shown in Fig. 41B represents a "1" bit. Each bit cell 1208 (64 image pixels) thus represents a single bit of the embedded data. Each sub-block 1206 includes 64 bit cells, and thus conveys 64 bits of embedded data.

- 131 -

(The nature of the image changes effected by the encoding follows the techniques set forth above under the heading MORE ON PERCEPTUALLY ADAPTIVE SIGNING; that discussion is not repeated here.)

In the illustrated embodiment, the embedded data includes two parts: control bits and
5 message bits. The 16 bit cells 1208A in the center of each sub-block 1206 serve to convey 16 control bits. The surrounding 48 bit cells 1208B serve to convey 48 message bits. This 64-bit chunk of data is encoded in each of the sub-blocks 1206, and is repeated 64 times in each signature block 1204.

A digression: in addition to encoding of the image to redundantly embed the 64
10 control/message bits therein, the values of individual pixels are additionally adjusted to effect encoding of subliminal gratitudes through the image. In this embodiment, the gratitudes discussed in conjunction with Fig. 29A are used, resulting in an imperceptible texturing of the image. When the image is to be decoded, the image is transformed into the spatial domain, the Fourier-Mellin technique is applied to match the graticule energy points with their expected positions, and
15 the processed data is then inverse-transformed, providing a registered image ready for decoding. (The sequence of first tweaking the image to effect encoding of the subliminal gratitudes, or first tweaking the image to effect encoding of the embedded data, is not believed to be critical. As presently practiced, the local gain factors (discussed above) are computed; then the data is encoded; then the subliminal graticule encoding is performed. (Both of these encoding steps
20 make use of the local gain factors.))

Returning to the data format, once the encoded image has been thus registered, the locations of the control bits in sub-block 1206 are known. The image is then analyzed, in the aggregate (i.e. considering the "northwestern-most" sub-block 1206 from each signature block 1204), to determine the value of control bit #1 (represented in sub-block 1206 by bit cell
25 1208Aa). If this value is determined (e.g. by statistical techniques of the sort detailed above) to be a "1," this indicates that the format of the embedded data conforms to the standard detailed herein (the Digimarc Beta Data Format). According to this standard, control bit #2 (represented by bit cells 1208Ab) is a flag indicating whether the image is copyrighted. Control bit #3

- 132 -

(represented by bit cells 1208Ac) is a flag indicating whether the image is unsuitable for viewing by children. Certain of the remaining bits are used for error detection/correction purposes.

— The 48 message bits of each sub block 1206 can be put to any use; they are not specified in this format. One possible use is to define a numeric "owner" field and a numeric "image/item" field (e.g. 24 bits each).

If this data format is used, each sub-block 1206 contains the entire control/message data, so same is repeated 64 times within each signature block of the image.

If control bit #1 is not a "1," then the format of the embedded data does not conform to the above described standard. In this case, the reading software analyzes the image data to determine the value of control bit #4. If this bit is set (i.e. equal to "1"), this signifies an embedded ASCII message. The reading software then examines control bits #5 and #6 to determine the length of the embedded ASCII message.

If control bits #5 and #6 both are "0," this indicates the ASCII message is 6 characters in length. In this case, the 48 bit cells 1208B surrounding the control bits 1208A are interpreted as six ASCII characters (8 bits each). Again, each sub-block 1206 contains the entire control/message data, so same is repeated 64 times within each signature block 1204 of the image.

If control bit #5 is "0" and control bit #6 is "1," this indicates the embedded ASCII message is 14 characters in length. In this case, the 48 bit cells 1208B surrounding the control bits 1208A are interpreted as the first six ASCII characters. The 64 bit cells 1208 of the immediately-adjointing sub-block 1220 are interpreted as the final eight ASCII characters.

Note that in this arrangement, the bit-cells 1208 in the center of sub-block 1220 are not interpreted as control bits. Instead, the entire sub-block serves to convey additional message bits. In this case there is just one group of control bits for two sub-blocks.

- 133 -

Also note than in this arrangement, pairs of sub-blocks 1206 contains the entire control/message data. so same is repeated 32 times within each signature block 1204 of the image.

Likewise if control bit #5 is "1" and control bit #6 is "0." This indicates the embedded
5 ASCII message is 30 characters in length. In this case, 2 x 2 arrays of sub-blocks are used for each representation of the data. The 48 bit cells 1208B surrounding control bits 1208A are interpreted as the first six ASCII characters. The 64 bit cells of each of adjoining block 1220 are interpreted as representing the next 8 additional characters. The 64 bits cells of sub-block 1222 are interpreted as representing the next 8 characters. And the 64 bit cells of sub-block 1224 are
10 interpreted as representing the final 8 characters. In this case, there is just one group of control bits for four sub-blocks. And the control/message data is repeated 16 times within each signature block 1204 of the image.

If control bits #5 and #6 are both "1"s, this indicates an ASCII message of programmable length. In this case, the reading software examines the first 16 bit cells 1208B surrounding the
15 control bits. Instead of interpreting these bit cells as message bits, they are interpreted as additional control bits (the opposite of the case described above, where bit cells normally used to represent control bits represented message bits instead). In particular, the reading software interprets these 16 bits as representing, in binary, the length of the ASCII message. An algorithm is then applied to this data (matching a similar algorithm used during the encoding process) to
20 establish a corresponding tiling pattern (i.e. to specify which sub-blocks convey which bits of the ASCII message, and which convey control bits.)

In this programmable-length ASCII message case, control bits are desirably repeated several times within a single representation of the message so that, e.g., there is one set of control bits for approximately every 24 ASCII characters. To increase packing efficiency, the tiling
25 algorithm can allocate (divide) a sub-block so that some of its bit-cells are used for a first representation of the message, and others are used for another representation of the message.

- 134 -

Reference was earlier made to beginning the decoding of the registered image by considering the "northwestern-most" sub-block 1206 in each signature block 1204. This bears elaboration.

Depending on the data format used, some of the sub-blocks 1206 in each signature block 1204 may not include control bits. Accordingly, the decoding software desirably determines the data format by first examining the "northwestern-most" sub-block 1206 in each signature block 1204; the 16 bits cells in the centers of these sub-blocks will reliably represent control bits. Based on the value(s) of one or more of these bits (e.g. the Digimarc Beta Data Format bit), the decoding software can identify all other locations throughout each signature block 1204 where the control bits are also encoded (e.g. at the center of each of the 64 sub-blocks 1206 comprising a signature block 1204), and can use the larger statistical base of data thereby provided to extract the remaining control bits from the image (and to confirm, if desired, the earlier control bit(s) determination). After all control bits have thereby been discerned, the decoding software determines (from the control bits) the mapping of message bits to bit cells throughout the image.

To reduce the likelihood of visual artifacts, the numbering of bit cells within sub-blocks is alternated in a checkerboard-like fashion. That is, the "northwestern-most" bit cell in the "northwestern-most" sub-block is numbered "0." Numbering increases left to right, and successively through the rows, up to bit cell 63. Each sub-block diametrically adjoining one of its corners (i.e. sub-block 1224) has the same ordering of bit cells. But sub-blocks adjoining its edges (i.e. sub-blocks 1220 and 1222) have the opposite numbering. That is, the "northwestern-most" bit cell in sub-blocks 1220 and 1222 is numbered "63." Numbering decreases left to right, and successively through the rows, down to 0. Likewise throughout each signature block 1204.

In a variant of the Digimarc beta format, a pair of sub-blocks is used for each representation of the data, providing 128 bit cells. The center 16 bit cells 1208 in the first sub-block 1206 are used to represent control bits. The 48 remaining bit cells in that sub-block, together with all 64 bit cells 1208 in the adjoining sub-block 1220, are used to provide a 112-bit message field. Likewise for every pair of sub-blocks throughout each signature block 1204. In such an arrangement, each signature block 1204 thus includes 32 complete representations of the

- 135 -

encoded data (as opposed to 64 representations in the earlier-described standard). This additional length allows encoding of longer data strings, such as a numeric IP address (e.g. URL).

Obviously, numerous alternative data formats can be designed. The particular format used can be indicated to the decoding software by values of one or more control bits in the
5 encoded image.

Other Applications

Before concluding, it may be instructive to review some of the other fields where principles of applicant's technology can be employed.

15 One is smart business cards, wherein a business card is provided with a photograph having unobtrusive, machine-readable contact data embedded therein. (The same function can be achieved by changing the surface microtopology of the card to embed the data therein.)

Another promising application is in content regulation. Television signals, images on the internet, and other content sources (audio, image, video, etc.) can have data indicating their
20 "appropriateness" (i.e. their rating for sex, violence, suitability for children, etc.) actually embedded in the content itself rather than externally associated therewith. Television receivers, internet surfing software, etc., can discern such appropriateness ratings (e.g. by use of universal

- 136 -

code decoding) and can take appropriate action (e.g. not permitting viewing of an image or video, or play-back of an audio source).

In a simple embodiment of the foregoing, the embedded data can have one or more "flag" bits, as discussed earlier. One flag bit can signify "inappropriate for children." (Others can be, e.g., "this image is copyrighted" or "this image is in the public domain.") Such flag bits can be in a field of control bits distinct from an embedded message, or can -- themselves -- be the message. By examining the state of these flag bits, the decoder software can quickly apprise the user of various attributes of the image.

(As discussed earlier, control bits can be encoded in known locations in the image -- known relative to the subliminal graticules -- and can indicate the format of the embedded data (e.g. its length, its type, etc.) As such, these control bits are analogous to data sometimes conveyed in prior art file headers, but in this case they are embedded within an image, instead of prepended to a file.)

The field of merchandise marking is generally well served by familiar bar codes and universal product codes. However, in certain applications, such bar codes are undesirable (e.g. for aesthetic considerations, or where security is a concern). In such applications, applicant's technology may be used to mark merchandise, either through in innocuous carrier (e.g. a photograph associated with the product), or by encoding the microtopology of the merchandise's surface, or a label thereon.

There are applications -- too numerous to detail -- in which steganography can advantageously be combined with encryption and/or digital signature technology to provide enhanced security.

Medical records appear to be an area in which authentication is important. Steganographic principles -- applied either to film-based records or to the microtopology of documents -- can be employed to provide some protection against tampering.

- 137 -

Many industries, e.g. automobile and airline, rely on tags to mark critical parts. Such tags, however, are easily removed, and can often be counterfeited. In applications wherein better security is desired, industrial parts can be steganographically marked to provide an inconspicuous identification/authentication tag.

5 In various of the applications reviewed in this specification, different messages can be steganographically conveyed by different regions of an image (e.g. different regions of an image can provide different internet URLs, or different regions of a photocollage can identify different photographers). Likewise with other media (e.g. sound).

10 Some software visionaries look to the day when data blobs will roam the datawaves and interact with other data blobs. In such an era, it will be necessary for such blobs to have robust and incorruptible ways of identifying themselves. Steganographic techniques again hold much promise here.

15 Finally, message changing codes — recursive systems in which steganographically encoded messages actually change underlying steganographic code patterns — offer new levels of sophistication and security. Such message changing codes are particularly well suited to applications such as plastic cash cards where time-changing elements are important to enhance security.

20 Again, while applicant prefers the particular forms of steganographic encoding detailed above, the diverse applications disclosed in this specification can largely be practiced with other steganographic marking techniques, many of which are known in the prior art. Likewise, while the specification has focused on applications of this technology to images, the principles thereof are generally equally applicable to embedding such information in audio, physical media, or any other carrier of information.

25 Finally, while the specification has been illustrated with particular embodiments, it will be recognized that elements, components and steps from these embodiments can be recombined in

different arrangements to serve different needs and applications, as will be readily apparent to those of ordinary skill in the art.

5 **POINTER TO A HEADER, AND/OR, POINTER TO A PROCESS**

 The basic philosophy surrounding the use of a "header" in a digital file has already been discussed, along with the idea that steganographically encoded information can serve many of these functions independently or in parallel with a traditional header.

10 This section attempts to further clarify this notion as well as further clarify the notion that attached information can act as a "pointer" to a remote database containing the information that might be otherwise contained in a header. Refer to the sections on header verification, the "bodier", and network linking methods and hot objects in general. This section also expands upon the idea that header information and other

15 kinds of attached information can be more than simply passive descriptive information, but can be fully executable code or other forms of "active information and network links."

 As previously developed in the section on hot objects, given that realtime connection of a desktop computer, a television set or a stereo system onto a network is becoming more and more commonplace, the absolute need for an attached header replete with full

20 information about a data object is diminishing, and the possibility of simply pointing to a remote yet connected database is becoming more and more practical. The persistent and immediate connection to the information is more fundamental than strictly needing the full information attached to the data object (where time to change information

25 presents a particular problem for object-attached information).

 Referring back to Fig. 16, and harkening back to the discussion on networked linking methods, we find the situation depicted in Fig. 60. The header information "joe's image" is now only located on some accessible remote database, and the information attached to the object in either the actual header, or steganographically

30 embedded, or both, now contains simply an address to the remote database location. Software applications would then be manufactured such that a networked connection would be assumed to be either active or capable of becoming so rather swiftly, and that

when a data object such as an image or a video clip was read in, the embedded address of the header information would be read, a quick retrieve of the database information would be performed, and the remote header information would be used just as if it were locally attached to the data object.

5 It is important to note that the steganographic embedded information does not necessarily need to be present in order to make such a system function, and vice versa, the traditional header need not be present if steganographic information is present. The two together represent a value-added redundancy and greater assurance that the link between the remote information and the object remain intact. Likewise, even if there
10 does exist some "master remote header" for a given data object, an "static-in-time copy" (by definition) of that master header file can still travel with the data object for the purposes of speed or for the purposes of servicing non-networked sties or applications, with reference back to the master if need be for any latest updates. Likewise, though the principles of this application focus largely on audio-visual signals and data objects,
15 many of these design principles apply to text files, software programs, basic data records, and the like, where steganography per se may not even be applicable. Likewise, any other kind of information can certainly still be attached to an object in addition to the above mentioned "header" information. Another way to look at it is that this new "header address" can simply be appended to or incorporated into any existing header
20 format such as JPEG, TIFF, etc.

It is also important to note that the information stored remotely need not simply be "passive information" such as a person's name or the latest price for some peddled item. It can also be computer instructions or software executable code that is associated with the data object. Thus a software application can read in a data object, reading at
25 the same time the remote database address, retrieve via a network the remote connection, then either display the remote information along with the data object information itself, or actually launch some local application much like JAVA based systems currently function, or the remote database can actually ship executable code which may then be invoked locally on the machine where the data object has been read.
30 In this fashion, a data object becomes a first stepping stone into some computing experience usually related to the content of the object itself, much as a cover of a book is a quick visual summary of the book's content where the challenge becomes enticing

third parties to enter.

The notion of a singular string of attached bits referring to a remote database wherein a master header is located we shall henceforth refer to as an "identity pointer". We shall refer to the database location containing the master information connected
5 with that data object as the "master identity."

As is well known even in the prior art of traditional mail, addresses can be hierarchical in nature, with ever finer levels of granularity in searching for the ultimate central location or master identity. So too with identity pointers, they can actually be a single string of hierarchical information parsed according to where an application or
10 machine needs to go in order to retrieve the full information associated with the data object and/or audio/visual signal. Thus, a sub-string inside the identity pointer might indicate, through proper programming of an application or instrument including address redirection tables, that the master identity for such and such object exists inside one's own computer at such and such directory and in such and such file. On the other hand,
15 a first section of the identity pointer might indicate a search is required on the world wide web for a particular rights management database or commerce server database, a second section of the identity pointer then "points" to an address at this particular database, wherein what is pointed to may be the ultimate location of the master identity, or the pointed-to location is actually a database of further pointers to a multitude of
20 databases, indicating that a third section of the identity pointer now has an address location of the ultimate master identity. This formula can clearly recurse further into fourth, fifth, and so on layers. The point is that it is not strictly required to have the identity pointer only point to some singular master database, but that there will in fact be a multitude of interconnected databases which will be the rule rather than the
25 exception. This includes databases which are more traditionally viewed as belonging in some kind of network which is not necessarily as large as the internet, such as local area networks, or even a personal computers series of hard drives, or whatever. The main overall point is that the identity pointer contains sufficient information in order to be able to find the master identity, where it is a given that an application or instrument can
30 read the master pointer, interpret it, and do the searching and retrieving of the master identity.

Those skilled in the art may notice that so far in this section no mention has been

made about encrypting either the identity pointer or the master identity or any communications involved in the overall process. Such improvements can certainly be made and enhance more of the "secure identification" aspects of an attached information system, or the related "verification" function that many commerce applications will be seeking. The basic principles of the invention can nevertheless be followed.

The identity pointer can also point to locations which are not (currently) viewed as strictly a database. For example, a data object can contain a sub-string which points to a particular application that needs to be launched in order to resolve the remainder of the identity pointer string. A more specific example of the general example might be that a sub-string in the identity pointer instructs a machine to launch a web browser, and the remainder of the string would be a traditional URL to some web page containing basic contact information to the creator of a video production. Users on a local machine might program the machine to their own liking in that they would "allow" content with these abilities to go ahead and launch browsers when instructed, or they can tell their local machine to ignore these requests by a given piece of content. Likewise, an sub-string in the identity pointer might instruct a multimedia database interface program to be launched and the location of the object in question sought, found, and displayed.

It can be readily seen that the general principles of the invention would indicate that applications and instruments which deal with data objects require less and less local intelligence regarding "interpreting and operating upon" attached information, yet instead they become better and better at being realtime connected to a broader network and know how to properly find and retrieve information from various locations on that broader network. Moreover, applications and instruments so enabled would also make use of well-known-in-the-art methods of launching new execution threads or entire new applications based on the retrieved information. The physical location of the master identity, as well as the "level" it exists on in the network hierarchy (from CPU RAM, through Operating System, through local storage disks, through local networks, up through global networks such as the Internet) becomes immaterial so long as an application or instrument is equipped with the realtime ability to find and retrieve the pertinent information.

We now turn toward a more detailed description of the drawings. Fig. 60 depicts

the essential concept behind the "master identity object" and a system built around it. ("Master Identity Object", or MIO, is here used as a convenience and should not be construed as limiting the kind of information pointed to by an embedded pointer; i.e., the MIO could easily point to executable code or executable sequences for example rather than identity information). Indeed, the system is essential to the definition of the object. Quite simply, a header is replaced by a pointer to a header, and the header information itself is located in some accessible location. The pointer 1690 represents whatever complete address information is needed in order for an instrument or application to locate the header information, 1694, located in some binary storage location (or database) 1692. The data stream 802 is whatever normal data that might be contained in a typical digital image file or binary record or software program or whatever. The pointer, 1690, can also be embedded steganographically as depicted, or alternatively, it can ONLY be associated with the data stream steganographically.

Fig. 61 begins to dissect Fig. 60, trying to depict a closer view of a practical MIO ("master identity object"). Reverence is paid to the idea that common data objects and common data formats will be used for some time to come, and in many ways the MIO information will be appended. Thus we start with a traditional header 1700 and the associated data stream 802. The MIO pointer is now split out into three pieces, a "tag" which identifies the following information as a pointer to a remote database. The pointer is now conveniently split into a pointer to a pointer 1704, effectively giving the address of some database in a sea of potential databases, while pointer 1706 is the ultimate address on the found database pointed to by 1704. Indeed, 1704 could also be split out, just as in surface mail type addressing.

Fig. 61 also indicates that multiple layers of MIO information can be added to an object. One example of this would be that an organization would track and account for all incoming and outgoing data objects and wish to provide a pointer to that given objects lifetime within the organizations data network, yet respect any previous MIO information that the creator or distributor of a data object already attached to that object. There are numerous other examples, including derivative work applications or situations where a data object evolves.

Fig. 62 takes a closer look at the connectivity components of an MIO system, looked at through the eyes of an instrument or application needing to retrieve the master

identity information for a given MIO-enabled data object. An arbitrary data object 1710 enters a highly genericized "node" 1712. Typically, this node might be an "open file" procedure in a software application, or it might be a "search for all image files" operation initiated in a computer's operating system, or it might be a digital receiver in a television set.

The node 1712 is meant to represent the place where a given data object enters a given instrument or application and is scanned for MIO traits. The node 1712's output is simply that the data object either is or is not an MIO type object. If it discovers an MIO object as in 1716, it must next ask the simple question of whether or not the given instruments cares or not, essentially asking whether or not to go fetch the information pointed to by the embedded pointer. If it is not configured to do so, it goes on with its work, 1718. If it is configured to go fetch the associated remote information (1720), it then simply goes through a basic address resolution process to locate the master identity. A first low level search might be typical wherein the "master identity" (which we have amply explained might actually be a search for further instructions rather than only "looking for a name") is potentially found on a user's own desktop computer or local instrument. This might be as low level as a call to some memory location in the CPU system 1724, perhaps retrieving and therefor instigating instructions to prepare for video information to follow. Or it might be a pointer 1726 to some location in a display device, perhaps as a precursor to understanding the current display settings that a user is operating under.

The need for multiple layers of MIO information as depicted in Fig. 61 becomes more apparent in these examples as well, in that the lower level retrieving will often be precursors to more advanced actions. Likewise, an address might be found in the operating system 1728, ROM, RAM, hard drives, etc. (1730). A token indication of a firewall 1732 is depicted showing that some addresses will only be accessible to a "local object" by design. Finally, perhaps the more common use of true "identity" will come from a route to a broader network 1734. Depicted is a certain remote database 1736 where the given data object might have its master identity. These two basic ideas of "local" identity and "remote" identity are depicted below as 1738 and 1740, but it should be appreciated that these can become quite blurry in their distinctions.

It should be understood that one object of this invention is that, for many

practical systems, the locating process of finding and retrieving the master identity information or instructions is meant to take place in realtime, that is, is typically less than a second or two ideally, and no more than ten seconds or twenty seconds even in case where there is in the multiples of kilobytes to retrieve. This kind of time performance is generally possible today on the internet even with sub-100Kbit per second modems, though less than a second is generally out of reach of current modem-based systems using standard phone lines and standard internet communications protocols.

Fig. 63 traces out a generic process flow within an instrument or application with a slightly different twist than depicted in Fig. 62. Again, any device capable of processing a data object reads said object in, 1750, and asks the simple question of whether or not remote associated information exists for that object 1752. If not, fine, 1754, go about your way. If so, it asks should I care, 1758, if not (1760, go about your way. If so, 1762, go fetch the information (or series of information if there is recursive MIO information). Box 1766 then explicitly shows that the user's device will still effect what happens even after the remote data is retrieved.

The permutations and combinations of operations from this point forward are as broad as general computing principles permit. As one tiny example, the user may have configured the device such that if the retrieved information does not match certain "key words" that the user has chosen, then don't even bother displaying an image (assuming that is the data object in question). Essentially, the retrieved information from 1764 and the user configuration 1766 set up a generally unlimited set of new instructions to the device. Another example might be that the user configuration 1766 says to go ahead and follow whatever instructions come back with the retrieved information, such as "launch a browser and follow an URL to some JAVA web site and don't mind sending out my credit card information if they ask". This is deliberately melodramatic to illustrate that the basic branching at 1770 is as broad as general computing principles permit.

Fig. 64 briefly illustrates that both embedded information as well as pointed-to information can actually sequences of instructions, including decision branched instructions, not merely singular instructions or singular blocks of passive descriptive information. This idea is encapsulated in the word "script" as is generally well known

in the art of computing.

Fig. 65 essentially depicts an MIO system operated in "reverse" from that of the previous few figures. In general, the idea here is that the master identity wants to know and collect all migration activities of its associated data object, and the system depicted in Fig. 65 is essentially a way to do this. Again, we see a generic data object 1820 entering some generic application or instrument 1822, here referred to as a "reporting node". The first question still is "are you an MIO?" 1824. 1826 alludes to the idea of reporting non-hits to some kind of master MIO system database 1830, a database which might track the broad traffic patterns of MIO object vs. Non-MIO objects. If the node finds an MIO, the next logical question is "may" I send a message to its master identity in order to inform the master identity that this object passed this location at this particular time? (1834). This is important since there will be many situations due to privacy and security considerations where such reporting would be intolerable 1834. However, there may also be many situations where this reporting information is valuable to all users of MIO based systems, and a "count" message will be sent to the master identity location, where locating the master identity follows previous descriptions. Details of how the master identity collects messages are left silent here as this form of data collection technology is readily available in today's marketplace and computing systems, where ultra-mundane methods such as sending a self-contained e-mail to some e-mail address might even suffice as a methodology of counting (horribly inefficient but indicative of the vast arrays of methods to have a master identity database keep track of hundreds if not millions of copies of "itself").

In view of the wide variety of implementations and applications to which the principles of this technology can be put, it should be apparent that the detailed embodiments are illustrative only and in no way limit the scope of my invention. I claim as my invention all such embodiments as come within the scope and spirit of the following claims and equivalents thereto.

CLAIMS

1. A multi-computer system including a network for embedding and reading a watermark, the system comprising:

5 a first digital electrical computer system comprising a first digital electrical computer electrically connected to a first input device, to a first output device, and to a first memory storing a plurality of creator identifiers and creator contact data corresponding to each of the creator identifiers;

10 a second digital electrical computer system comprising a second digital electrical computer electrically connected to a second input device and to a second output device, the second digital electrical computer being programmed to embed a watermark in a digital photographic image, the watermark including one of the plurality of creator identifiers;

15 a third digital electrical computer system comprising a third digital electrical computer electrically connected to a third input device and to a third output device, the third digital electrical computer being programmed to read the watermark in the digital photographic image to reveal the one of the plurality of creator identifiers;

20 a network for communicating the revealed one of the plurality of creator identifiers to the first digital electrical computer to obtain the creator contact data corresponding to the one of the plurality of creator identifiers from the memory.

2. The system of claim 1, wherein:

the second digital electrical computer is programmed to automatically detect for a watermark when an image is first examined by the second digital electrical computer.

3. The system of claim 1, wherein:

25 the second digital electrical computer is programmed to selectably detect for a watermark when an image is examined by the second digital electrical computer.

4. The system of claim 1, wherein:

the network includes the Internet;

the watermark includes information identifying a World Wide Web site; and

30 wherein

the third digital electrical computer system is programmed to load a World Wide Web browser and connect to the World Wide Web site in response to the revealed one

of the plurality of creator identifiers.

5. The system of claim 1, wherein:

the watermark includes extended data including at least one member from the group consisting of an organization identifier, a transaction identifier, and an item identifier.

6. In a television set, audio stereo system, computing system, or other system capable of receiving, storing, transmitting, or processing an audio or visual signal, an improvement comprising:

means for adding to an initial audio or visual signal an additional digital information associated with the initial signal,

the additional information containing, in whole or in part, an identification number which is uniquely assigned to the individual initial signal or copies thereof,

wherein the identification number represents an address to a location in a remote database wherein a prestored set of information associated with the initial signal is contained, and

means coupled to the system containing a processor responsive to basic operating instructions for enabling the reading of the additional information and enabling the realtime retrieving from the remote database a subset of the prestored information via a communications link with the remote site.

7. The system of claim 6, wherein the prestored set of information comprises identification and description information about the signal or an object contained in or represented by the signal.

8. The system of claim 6, wherein the prestored set of information comprises a set of machine instructions executable to perform a task relative to the initial signal or about an object contained in the initial signal.

9. The system of claim 6, wherein the prestored set of information comprises a menu of options or commands.

10. The system of claim 1, including means for steganographically embedding the additional information in the initial signal.

11. The system of claim 1, wherein there also exists a second number as part of the additional information, wherein the second number is the address or identifier of the database itself to which the first identification points inside, in the context of

multiple databases serving a multitude of audio and visual signals.

12. The system of claim 4, first and second numbers steganographically embedded in the initial signal.

13. The system of claim 1, wherein there is additionally an instrument or receiving station which monitors a multiple number of audio/visual public channels searching and recording the identification numbers and reporting sitings.

14. In an image processing method that includes steganographically decoding an input two-dimensional image to extract a multi-bit code therein, the image comprising a two-dimensional array of pixels, an improvement comprising:

10 transforming the image into the spatial frequency domain;

pattern matching the transformed image so spatial frequencies obtained by said transforming step coincide with reference spatial frequencies, to thereby effect registration of the transformed image;

inverse-transforming the transformed image to yield a registered image;

15 identifying, in the registered image, a plurality of regions that encode a first control bit, said regions being distributed through the registered image in a regular array;

performing a statistical analysis over at least said plurality of regions to determine whether the first control bit has first or second values;

20 if said control bit has the first value, performing a first decoding process on the image to extract the code therefrom; and

if said control bit has the second value, performing a second decoding process on the image to extract the code therefrom, the second decoding process being different than the first.

FIG. 1





FIG. 4



2 /55

FIG. 2

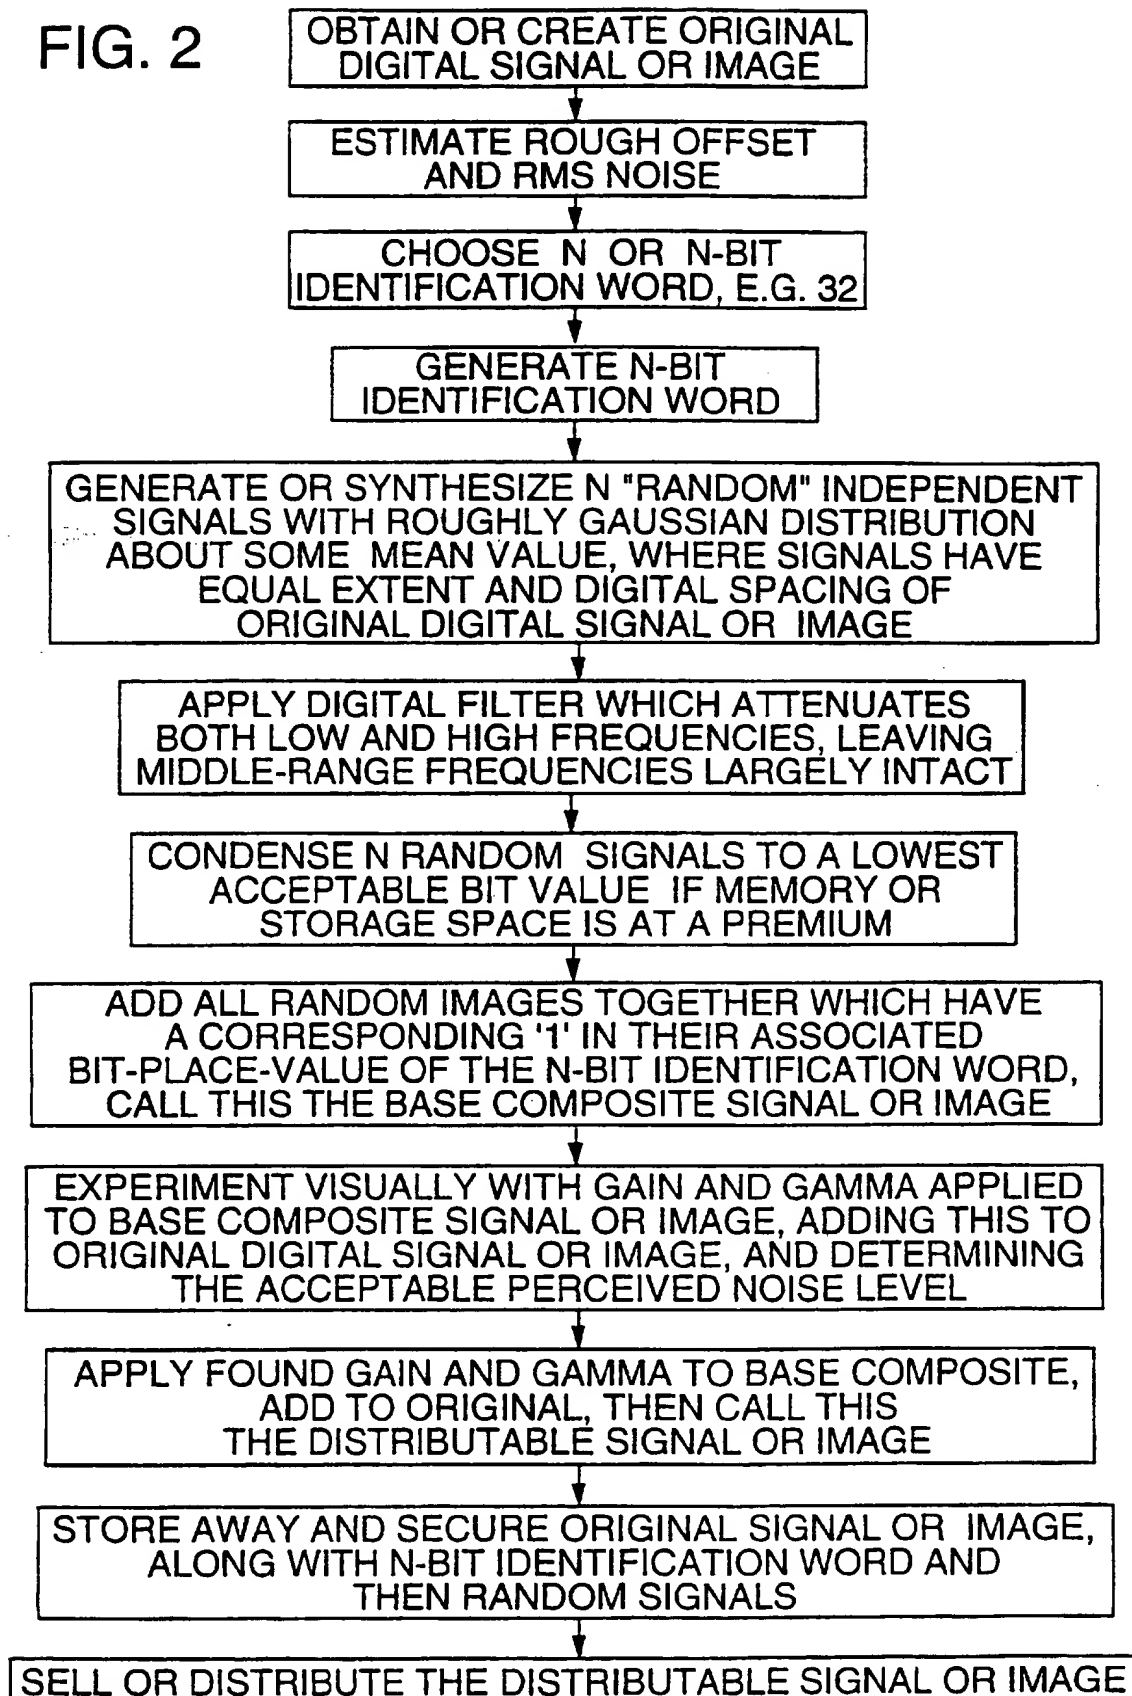


FIG. 3

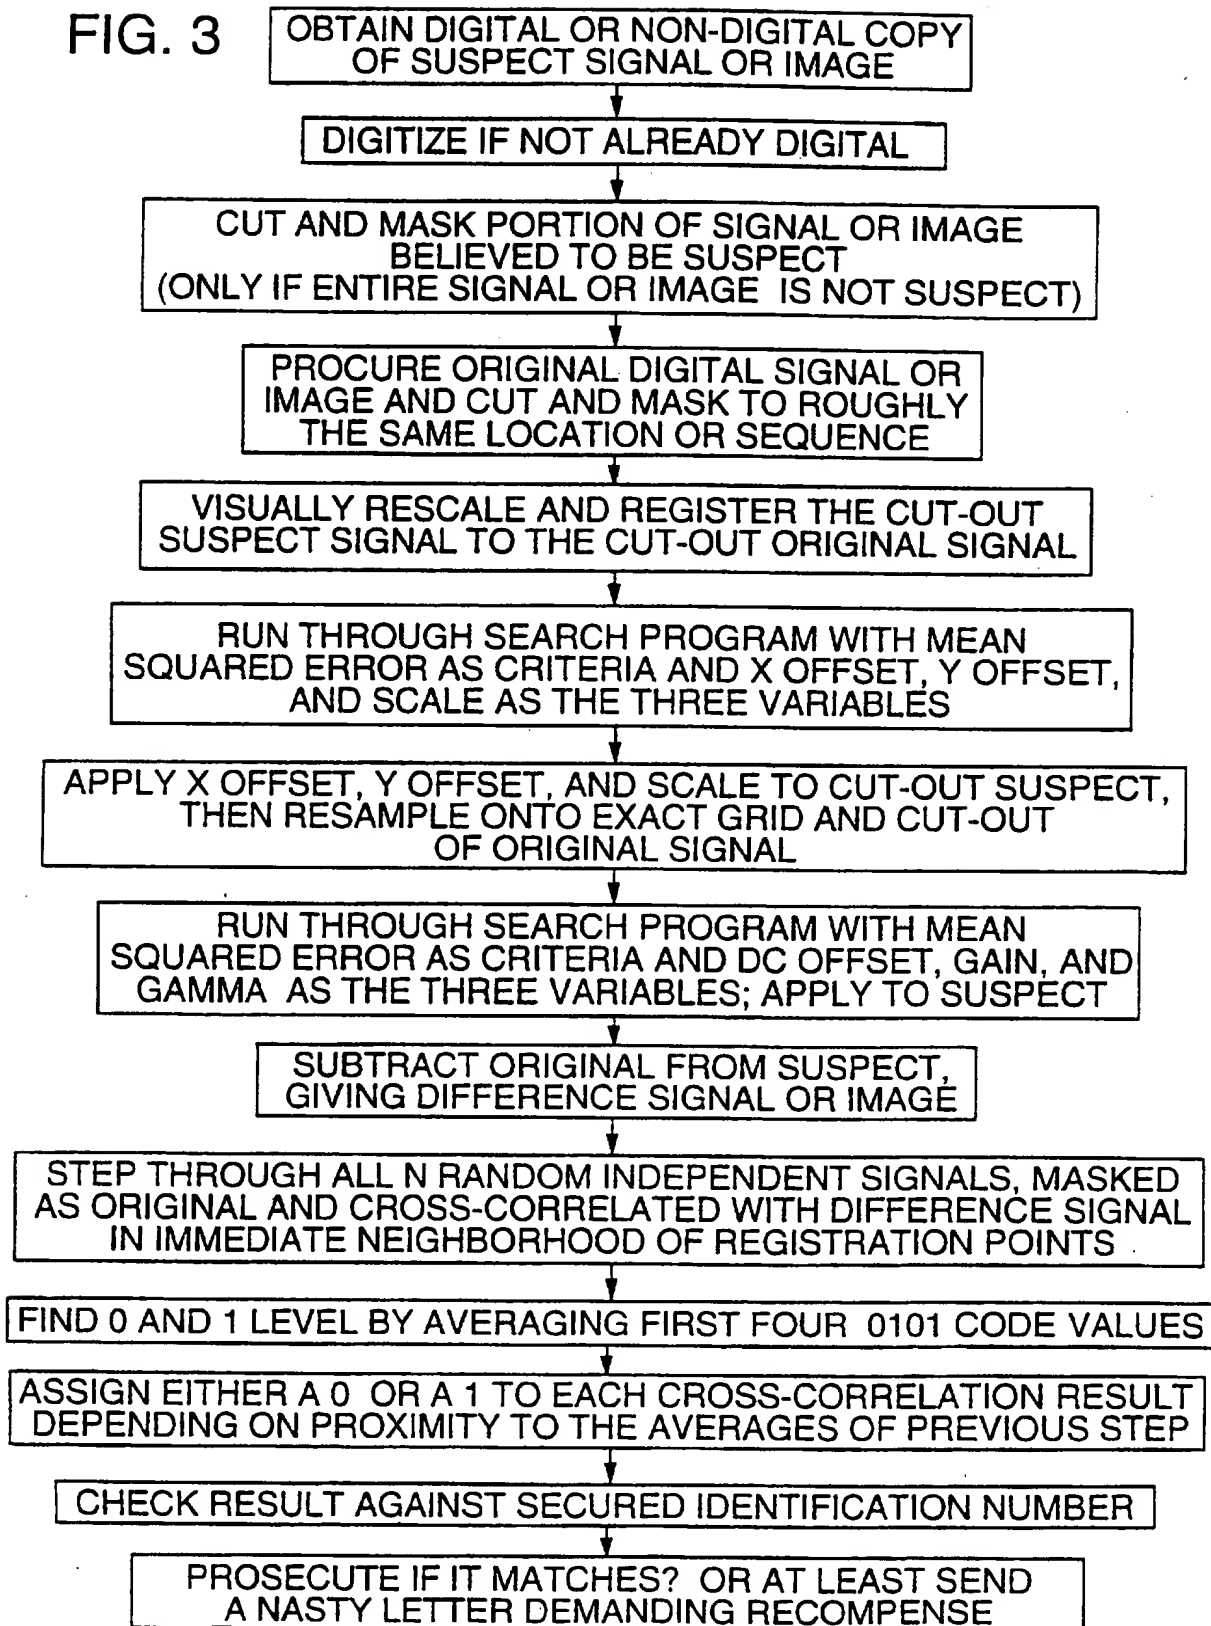


FIG. 5





FIG. 6



5 / 55





FIG. 9A




FIG. 9B




FIG. 9C

BORDER
CONTINUITY
404





FIG. 10



7 / 55

FIG. 11




FIG. 12

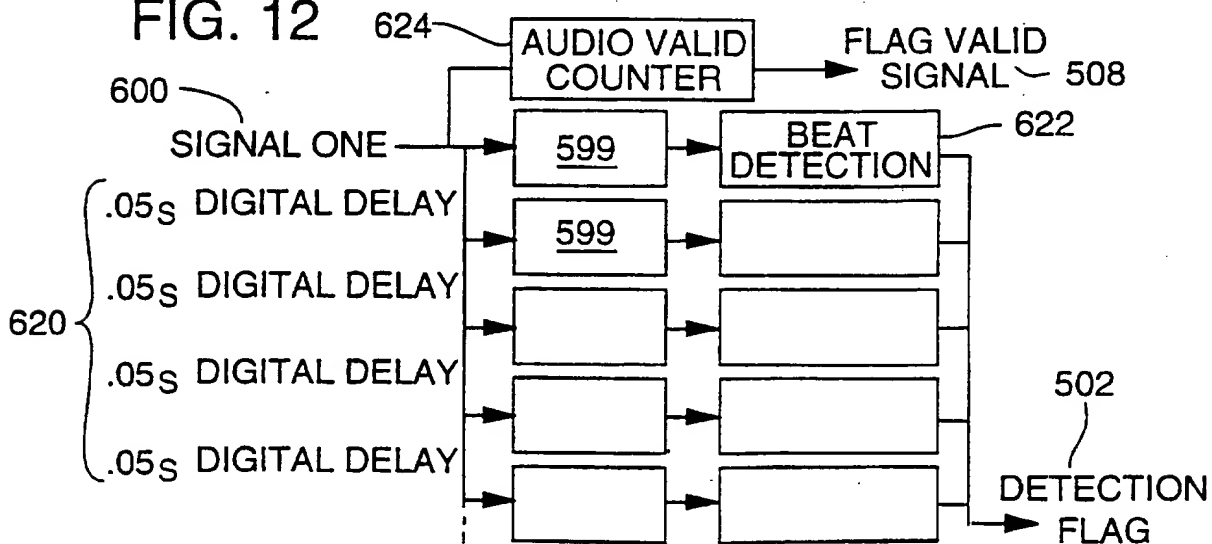


FIG. 13






FIG. 14




MEAN-REMOVED HISTOGRAMS OF
DIFFERENCE SIGNAL AND KNOWN EMBEDDED
CODE SIGNAL




MEAN-REMOVED HISTOGRAMS OF
FIRST DERIVATIVES (OR SCALAR GRADIENTS
IN CASE OF AN IMAGE)

10/55



11/55

FIG. 16



12/55

FIG. 17




FIG. 18






FIG. 19




QUEST FOR MOSAICED KNOT PATTERNS WHICH "COVER" AND
ARE COEXTENSIVE WITH ORIGINAL IMAGE;
ALL ELEMENTAL KNOT PATTERNS CAN CONVEY THE SAME
INFORMATION, SUCH AS A SIGNATURE, OR EACH CAN CONVEY A
NEW MESSAGE IN A STEGANOGRAPHIC SENSE

14/55




16/55



17/55


FIG. 23



CONTAINS RUDIMENTARY OPTICAL SCANNER,
MEMORY BUFFERS, COMMUNICATIONS DEVICES,
AND MICROPROCESSOR

CONSUMER MERELY PLACES CARD INTO WINDOW
AND CAN, AT THEIR PREARRANGED OPTION, EITHER
TYPE IN A PERSONAL IDENTIFICATION NUMBER
(PIN, FOR ADDED SECURITY) OR NOT. THE TRANSACTION
IS APPROVED OR DISAPPROVED WITHIN SECONDS.

FIG. 24




19/55

FIG. 25 TYPICAL TRANSACTION STEPS

1. READER SCANS IMAGE ON CARD, STORES IN MEMORY, EXTRACTS PERSON'S ID
2. OPTIONAL: USER KEYS IN PIN NUMBER
 3. READER CALLS CENTRAL ACCOUNT DATA NETWORK, HANDSHAKES
 4. READER SENDS ID, (PIN), MERCHANT INFORMATION, AND REQUESTED TRANSACTION AMOUNT TO CENTRAL NETWORK
 5. CENTRAL NETWORK VERIFIES ID, PIN, MERCHANT INFO, AND ACCOUNT BALANCE
 6. IF OK, CENTRAL NETWORK GENERATES TWENTY-FOUR SETS OF SIXTEEN DISTINCT RANDOM NUMBERS, WHERE THE RANDOM NUMBERS ARE INDEXES TO A SET OF 64K ORTHOGONAL SPATIAL PATTERNS
 7. CENTRAL NETWORK TRANSMITS FIRST OK, AND THE SETS OF RANDOM NUMBERS
8. READER STEPS THROUGH THE TWENTY-FOUR SETS
 - 8A. READER ADDS TOGETHER SET OF ORTHOGONAL PATTERNS
 - 8B. READER PERFORMS DOT PRODUCT OF RESULTANT PATTERN AND CARD SCAN, STORES RESULT
 9. READER TRANSMITS THE TWENTY-FOUR DOT PRODUCT RESULTS TO CENTRAL NETWORK
 10. CENTRAL NETWORK CHECKS RESULTS AGAINST MASTER
 11. CENTRAL NETWORK SENDS FINAL APPROVAL OR DENIAL
 12. CENTRAL NETWORK DEBITS MERCHANT ACCOUNT, CREDITS CARD ACCOUNT

20/55

FIG. 26
THE NEGLIGIBLE-FRAUD CASH CARD SYSTEM



A BASIC FOUNDATION OF THE CASH CARD SYSTEM IS A 24-HOUR INFORMATION NETWORK, WHERE BOTH THE STATIONS WHICH CREATE THE PHYSICAL CASH CARDS, 950, AND THE POINT-OF-SALES, 984, ARE ALL HOOKED UP TO THE SAME NETWORK CONTINUOUSLY

21/55





FIG. 28



22/55

FIG. 27A






FIG. 27B




23/55

FIG. 29



NON-HARMONIC SPATIAL FREQUENCIES ALONG THE 45 DEGREE AXES, GIVING RISE TO A WEAWE-LIKE CROSS-HATCHING PATTERN IN THE SPATIAL DOMAIN


FIG. 30



NON-HARMONIC CONCENTRIC CIRCLES IN UV PLANE, WHERE PHASE HOPS QUASI-RANDOMLY ALONG EACH CIRCLE, GIVING RISE TO PSEUDO RANDOM LOOKING PATTERNS IN THE SPATIAL DOMAIN


24/55

FIG. 29A




25/55

FIG. 31A




PHASE OF SPATIAL
FREQUENCIES ALONG
FORWARD 45 DEGREE
AXES, 1008

FIG. 31B




PHASE OF SPATIAL
FREQUENCIES ALONG
BACKWARD 45 DEGREE
AXES, 1010

FIG. 32A



PHASE OF SPATIAL
FREQUENCIES ALONG
FIRST CONCENTRIC RING,
1012

FIG. 32B



PHASE OF SPATIAL
FREQUENCIES ALONG
SECOND CONCENTRIC RING,
1014


FIG. 32C



PHASE OF SPATIAL
FREQUENCIES ALONG
THIRD CONCENTRIC RING,
1016


26/55

FIG. 33A




ANGLE A

FIG. 33B




ANGLE B

FIG. 33C




POWER PROFILE ALONG ANGLE A, AS NORMALIZED BY ITS OWN MOVING AVERAGE; ONLY A MINIMAL AMOUNT EXCEEDS THRESHOLD, GIVING A SMALL INTEGRATED VALUE

FIG. 33D



POWER PROFILE ALONG ANGLE B, AS NORMALIZED BY ITS OWN MOVING AVERAGE; THIS FINDS STRONG ENERGY ABOVE THE THRESHOLD

FIG. 33E



27/55

FIG. 34A




FIG. 34B




FIG. 34C




FIG. 34D






FIG. 34E



28/55

FIG. 35A




SCALE = A; ADD ALL POWER VALUES AT THE
"KNOWN" FREQUENCIES", 1042

FIG. 35B



SCALE = B; ADD ALL POWER VALUES AT THE
"KNOWN" FREQUENCIES", 1044


FIG. 35C



"SCALED-KERNEL" BASED MATCHED FILTER; PEAK IS
WHERE THE SCALE OF THE SUBLIMINAL GRID WAS
FOUND, 1046


29/55

FIG. 36A



ARBITRARY ORIGINAL IMAGE
IN WHICH SUBLIMINAL
GRATICULES MAY HAVE BEEN PLACED

FIG. 36B



"COLUMN SCAN"
IS APPLIED ALONG A
GIVEN ANGLE THROUGH
THE CENTER OF THE
IMAGE

COLUMN-
INTEGRATED
GREY
VALUES,
1054

FIG. 36C





FIG. 36D




MAGNITUDE OF FOURIER TRANSFORM OF SCAN DATA

FIG. 37


PROCESS STEPS

1. SCAN IN PHOTOGRAPH
2. 2D FFT
3. GENERATE 2D POWER SPECTRUM, FILTER WITH E.G. 3X3 BLURRING KERNEL
4. STEP ANGLES FROM 0 DEGREES THROUGH 90 (1/2 DEG)
5. GENERATE NORMALIZED VECTOR, WITH POWER VALUE AS NUMERATOR, AND MOVING AVERAGED POWER VALUE AS DENOMINATOR
6. INTEGRATE VALUES AS SOME THRESHOLD, GIVING A SINGLE INTEGRATED VALUE FOR THIS ANGLE
7. END STEP ON ANGLES
8. FIND TOP ONE OR TWO OR THREE "PEAKS" FROM THE ANGLES IN LOOP 4, THEN FOR EACH PEAK...
9. STEP SCALE FROM 25% TO 400%, STEP ~1.01
10. ADD THE NORMALIZED POWER VALUES CORRESPONDING TO THE 'N' SCALED FREQUENCIES OF STANDARD
11. KEEP TRACK OF HIGHEST VALUE IN LOOP
12. END LOOP 9 AND 8, DETERMINE HIGHEST VALUE
13. ROTATION AND SCALE NOW FOUND
14. PERFORM TRADITIONAL MATCHED FILTER TO FIND EXACT SPATIAL OFFSET
15. PERFORM ANY "FINE TUNING" TO PRECISELY DETERMINE ROTATION, SCALE, OFFSET

FIG. 38



32/55






FIG. 41B

| | | |
|---|---|---|
| + | - | + |
| - | + | - |

FIG. 41A

| | | |
|---|---|---|
| - | + | - |
| + | - | + |




FIG. 42

35/55




FIG. 43

36/55

MarcCentre

file:///C:/Data/Downloa



Welcome to MarcCentre™ Digimarc's Locator Service



MarcCentre is an on-line service which links your up-to-date contact information to every image you watermark. You provide the details you want consumers of your image to see. Each time someone reads one of your watermarked images, they can access this information about you and your work with a single mouse click.


Subscribe now to give each of your images a persistent identity, communicating your copyright, and bringing consumers directly to you. Take advantage of the special offers available to you now, for a limited time.

To find out more about how MarcCentre works and how you can get started taking advantage of this locator service, browse the MarcCentre overview page.

If you are already a subscriber, you can go to Member Services to view or update your contact details.

To look up image creator/distributor information, you can search by creator ID or by other, more general, criteria.

For the technically minded, find out what's going on behind the scenes.



[Digimarc Copyright Notice](#) | [About the Site](#)

FIG. 44

37/55

Subscription :

Page 1



Thank you for signing up with MarcCentre. There are two parts to the application form (1) your contact information and (2) Agencies or Representative contact information. Simply fill in the blanks on this page and click "Part 1 Complete" below.

Name (First/Last)*

Company

Address*

City*

State/Province*

Zip/Postal Code*

Country*

Phone Number*

Fax Number

E-mail Address

Your Home Page

You can control how much of your contact information is displayed when a potential customer finds your profile on MarcCentre. Use the selections below to tell us what you want people to see.

Phone Number

☒ Show☐ Don't show

Mailing Address

☒ Show☐ Don't show

Email Account

☒ Show☐ Don't showYes, I want to be on your mailing lists ☒

Please select your profession (one only).

☒ Photographer☐ Illustrator☐ Other

In order to communicate your primary area of focus or specialty to potential customers, please select one of the options from the list below -or- type in your own in the space next to "Custom Specialty"

Photographer

(none)

Illustrator

(none)

Custom Specialty

FIG. 45

38/55

Subscription :

Page 2 of 2

Digimarc displays an Image of the Day showcase on our web site. Would you like to participate?

☒ Yes.

☐ Not at this time.

For security purposes, you will need to setup a password that allows you to edit your profile information. This password along with your Creator ID are required before you can update your contact information.

Password* Re-enter to verify*

| | |
|--|--|
| | |
|--|--|

The standard fee for the MarcCentre service is \$150 (US). From now until December 31st, 1996 you can subscribe for only \$79 (US), over 50% off the regular price!

For security reasons, the following credit card information is captured using secure Digimarc displays an Image of the Day showcase on our web site. Would Credit you like to participate?

☒ Via: ☒ Yes.

☐ MasterCard

☐ American Express

Card Number*

Name Shown on Card*

Expiration
(MM/YY)*

| |
|--|
| |
|--|

| |
|--|
| |
|--|

| |
|--|
| |
|--|

When you are satisfied with your entries on this form press the "Part 1 Complete" button and continue with the second part of the subscription sign-up. If you wish to cancel, press the Back button on your browser.

Part 1 Complete

Digimarc Copyright Notice

FIG. 46

SUBSTITUTE SHEET (RULE 26)

39/55




FIG. 47

40/55

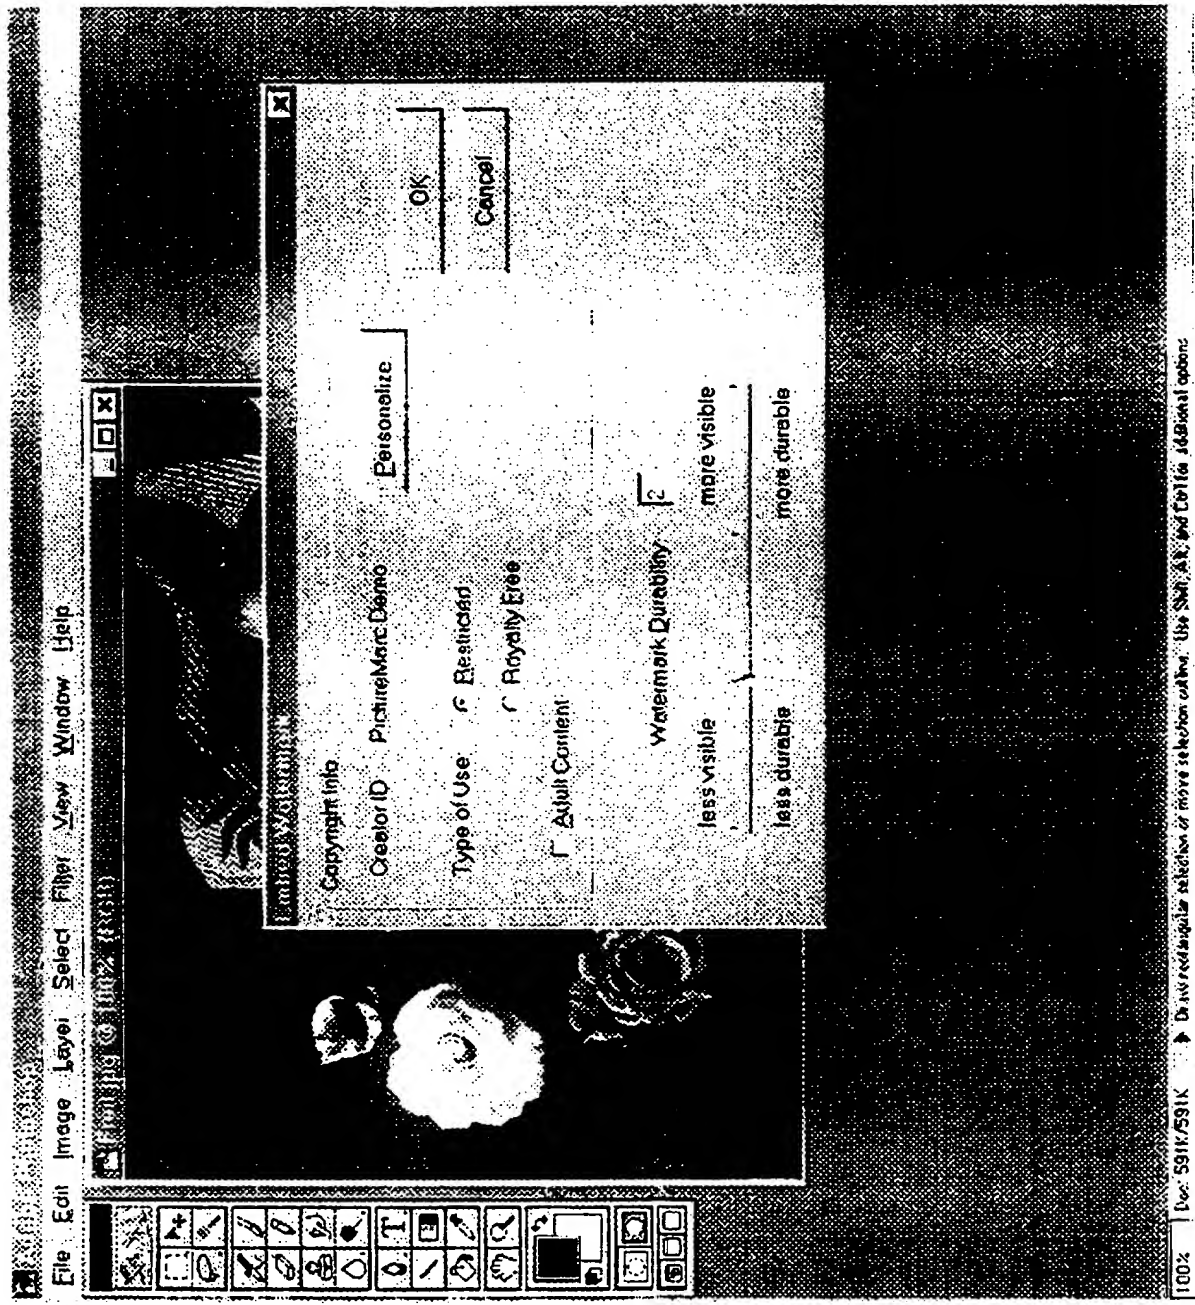


FIG. 48

41/55

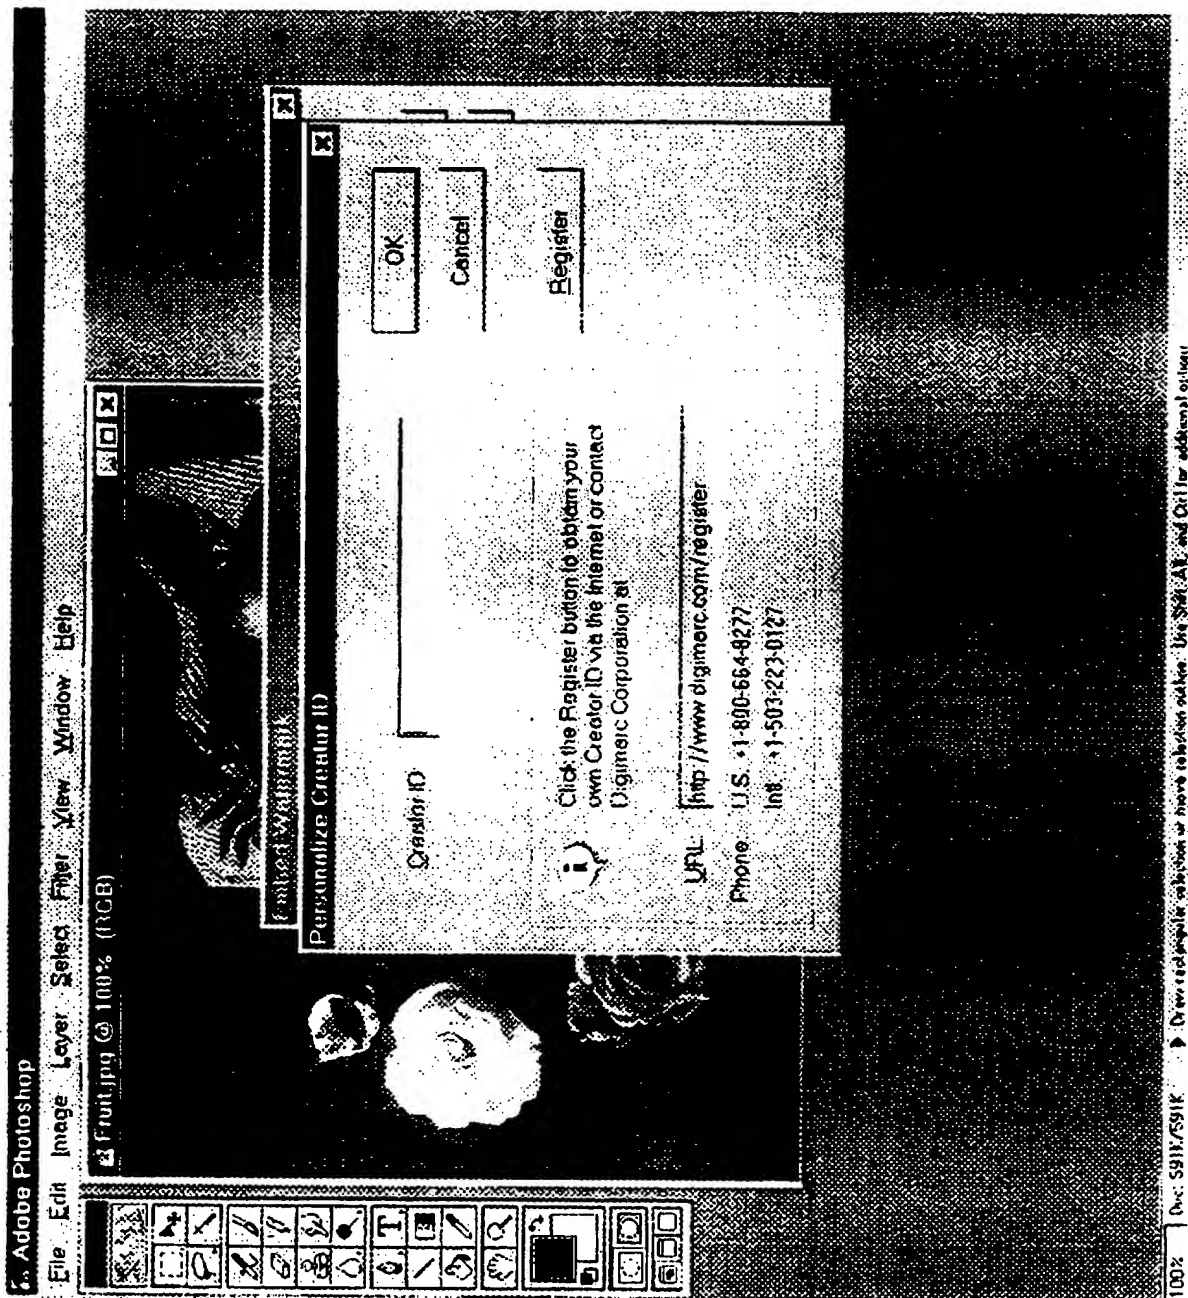


FIG. 49

42/55




FIG. 50

43/55



FIG. 51

44/55

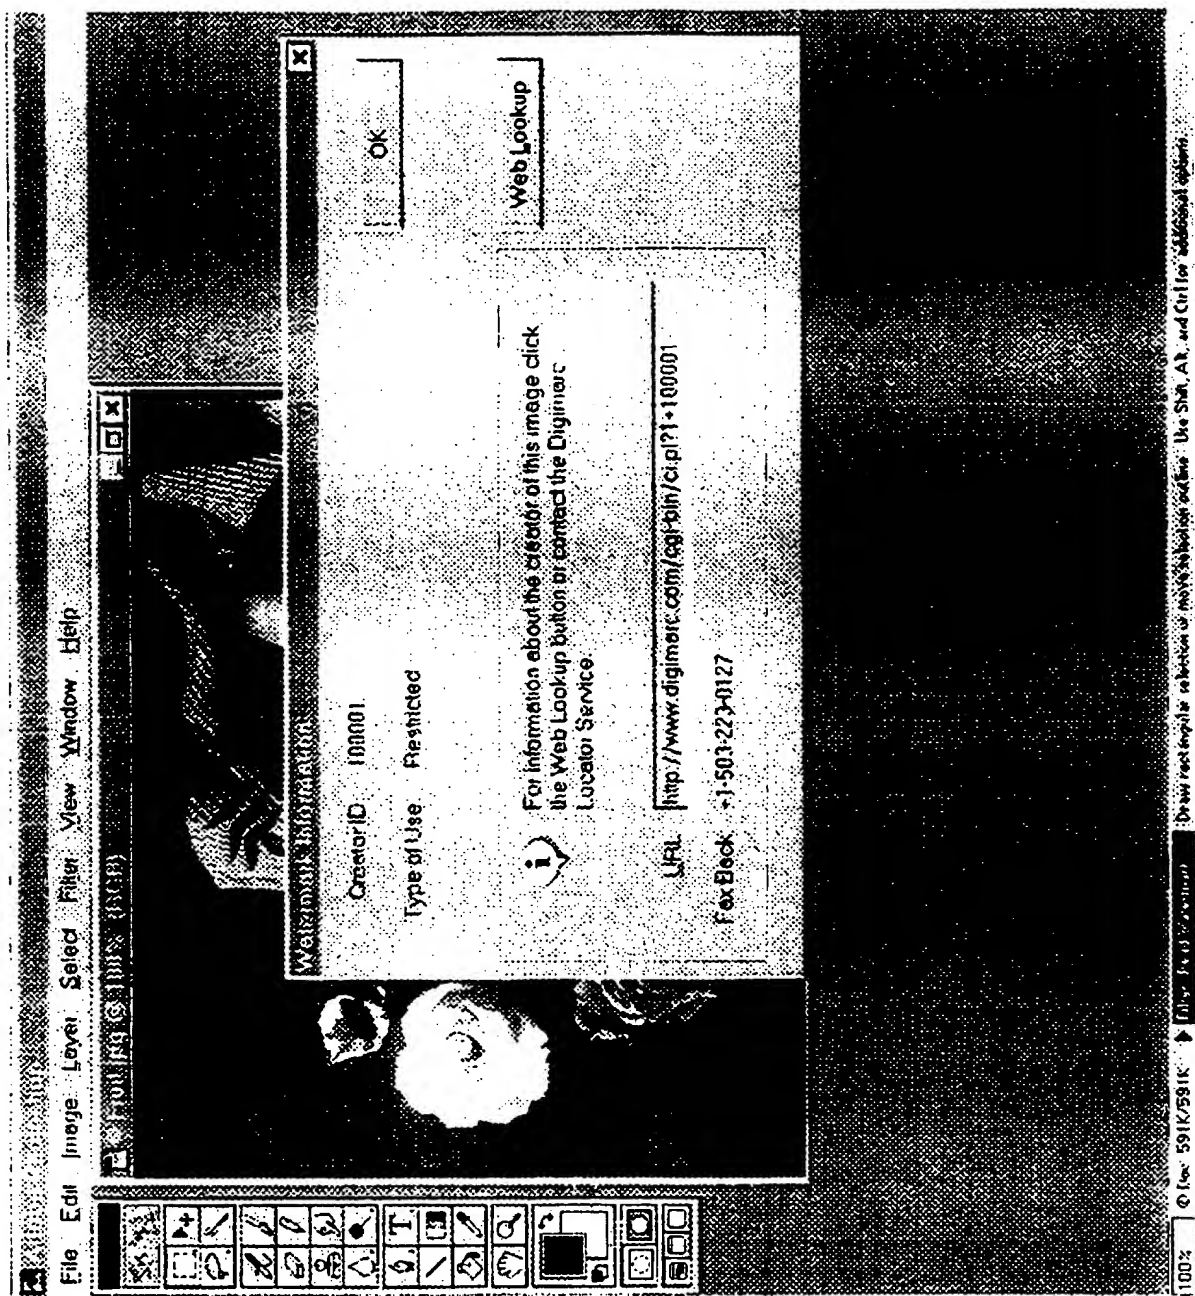


FIG. 52

45/55

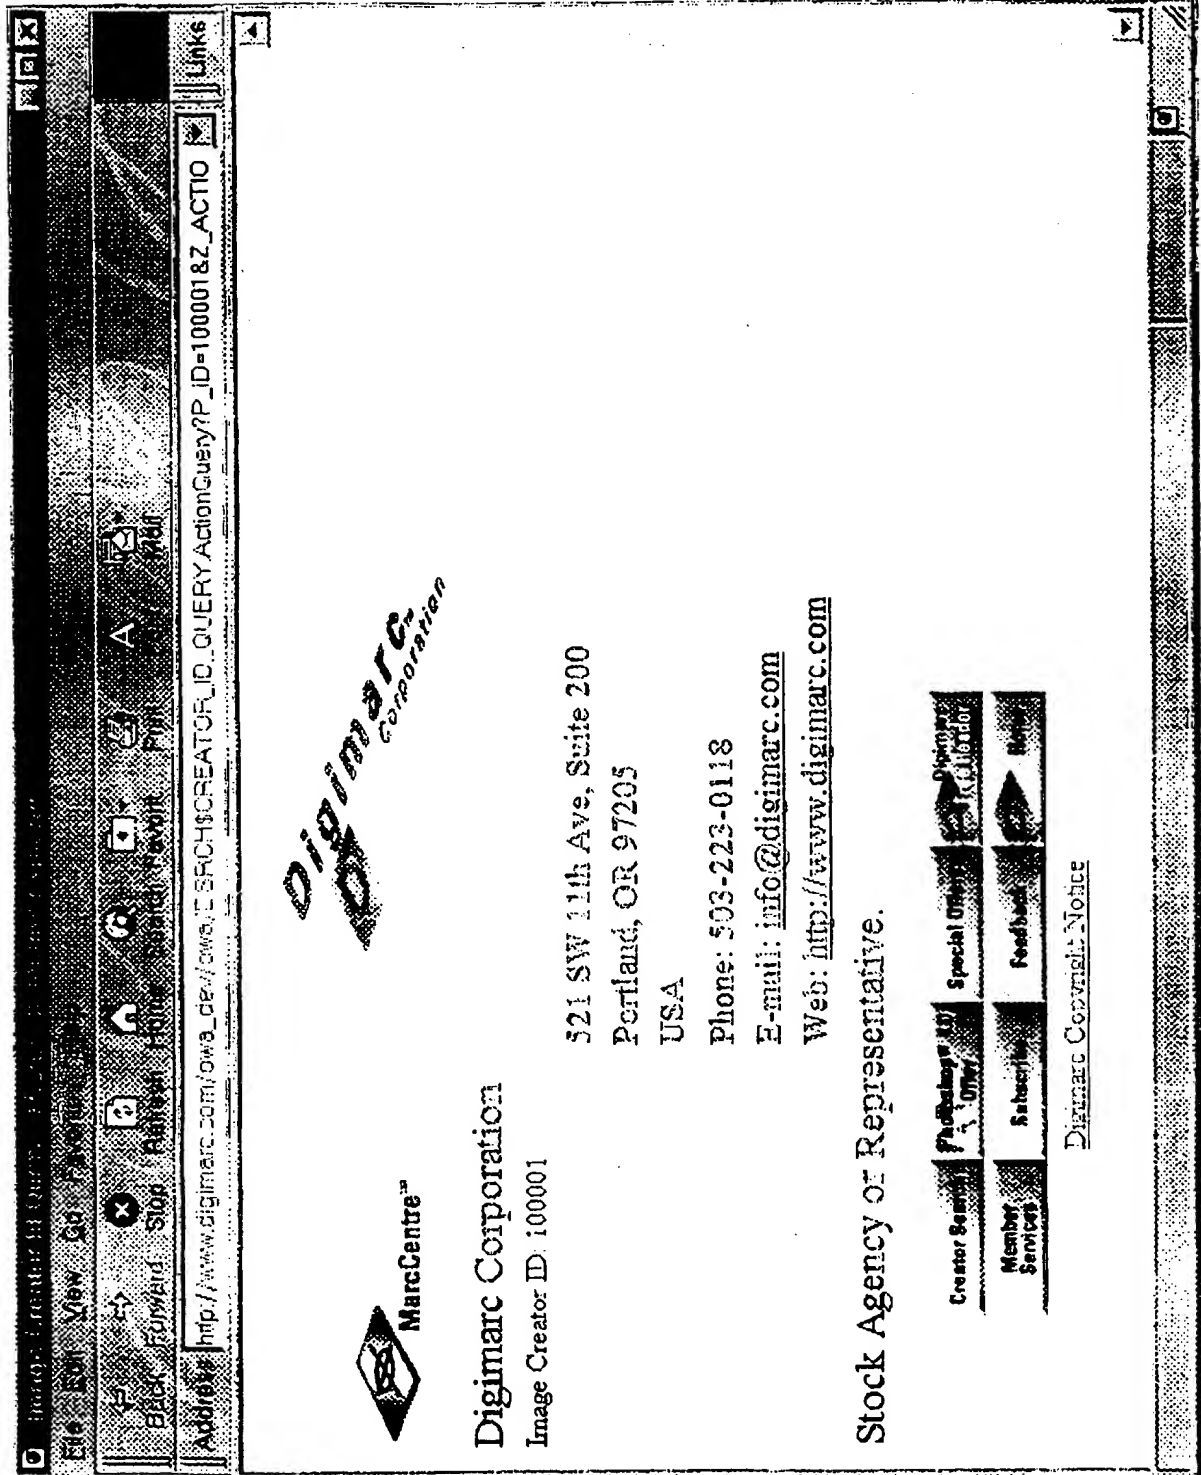



FIG. 53

Subscription :

Page 1 of 1




MarcCentre Image Creator Search

Enter the Image Creator ID in the box below and press "Submit Search". If it is a valid ID, the contact information details will be listed.

Image Creator ID:

If you want to search for specific image creators by a variety of criteria such as last name, specialty, city and/or state then press the "General Search" button below.



Digimarc Copyright Notice

FIG. 54

47/55

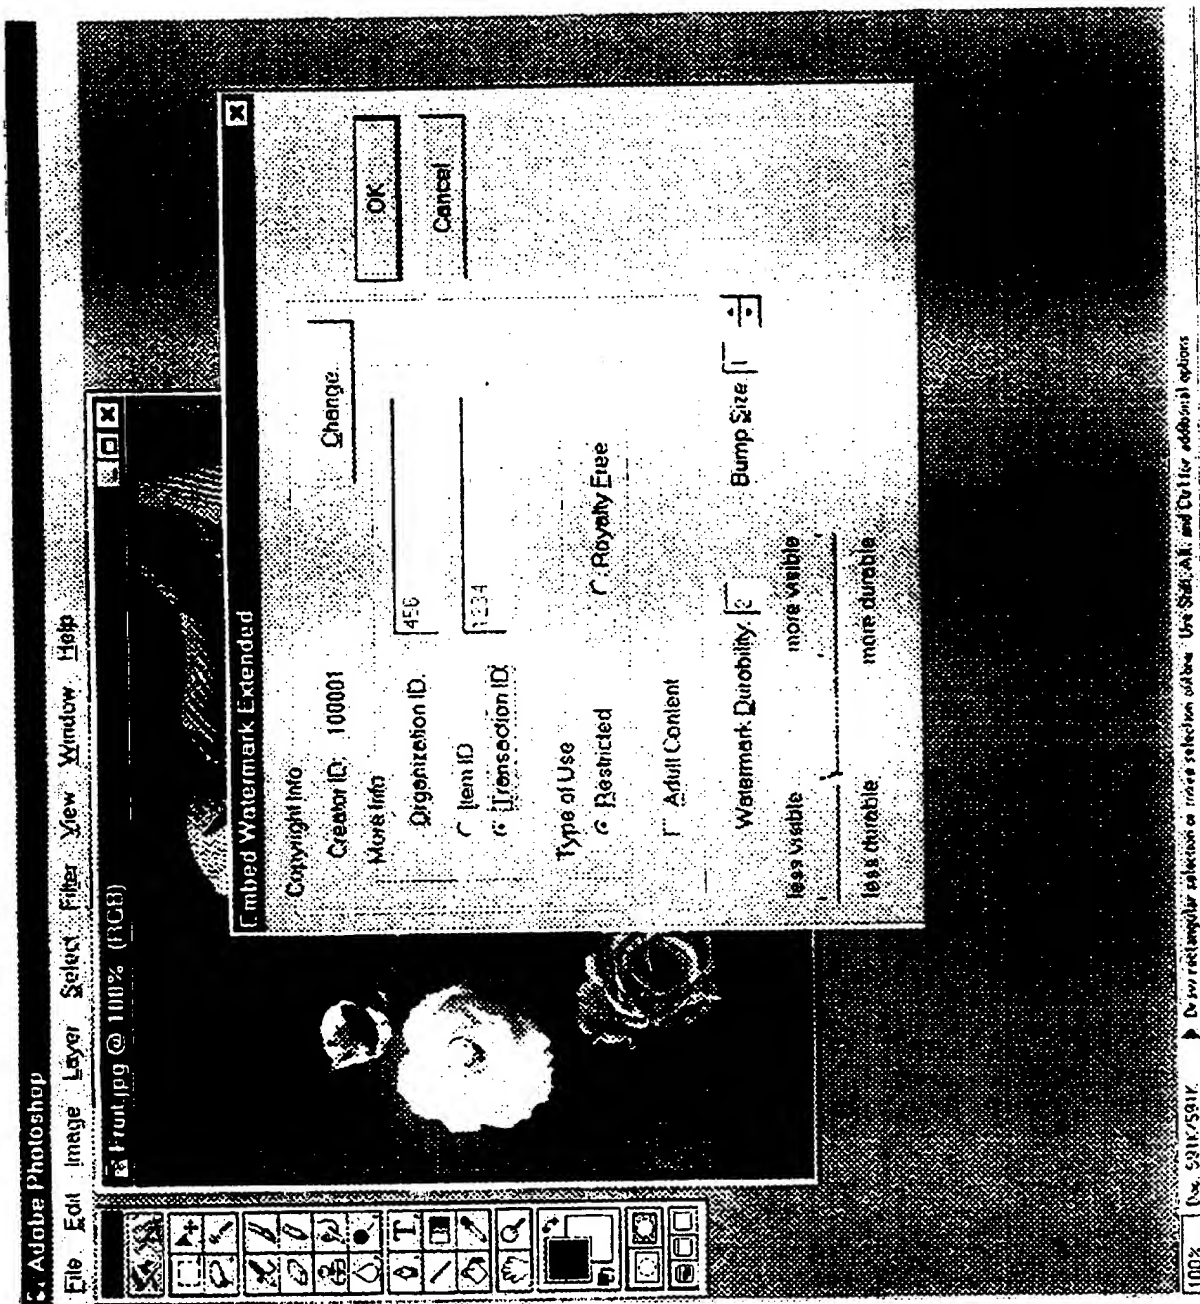


FIG. 55

48/55.

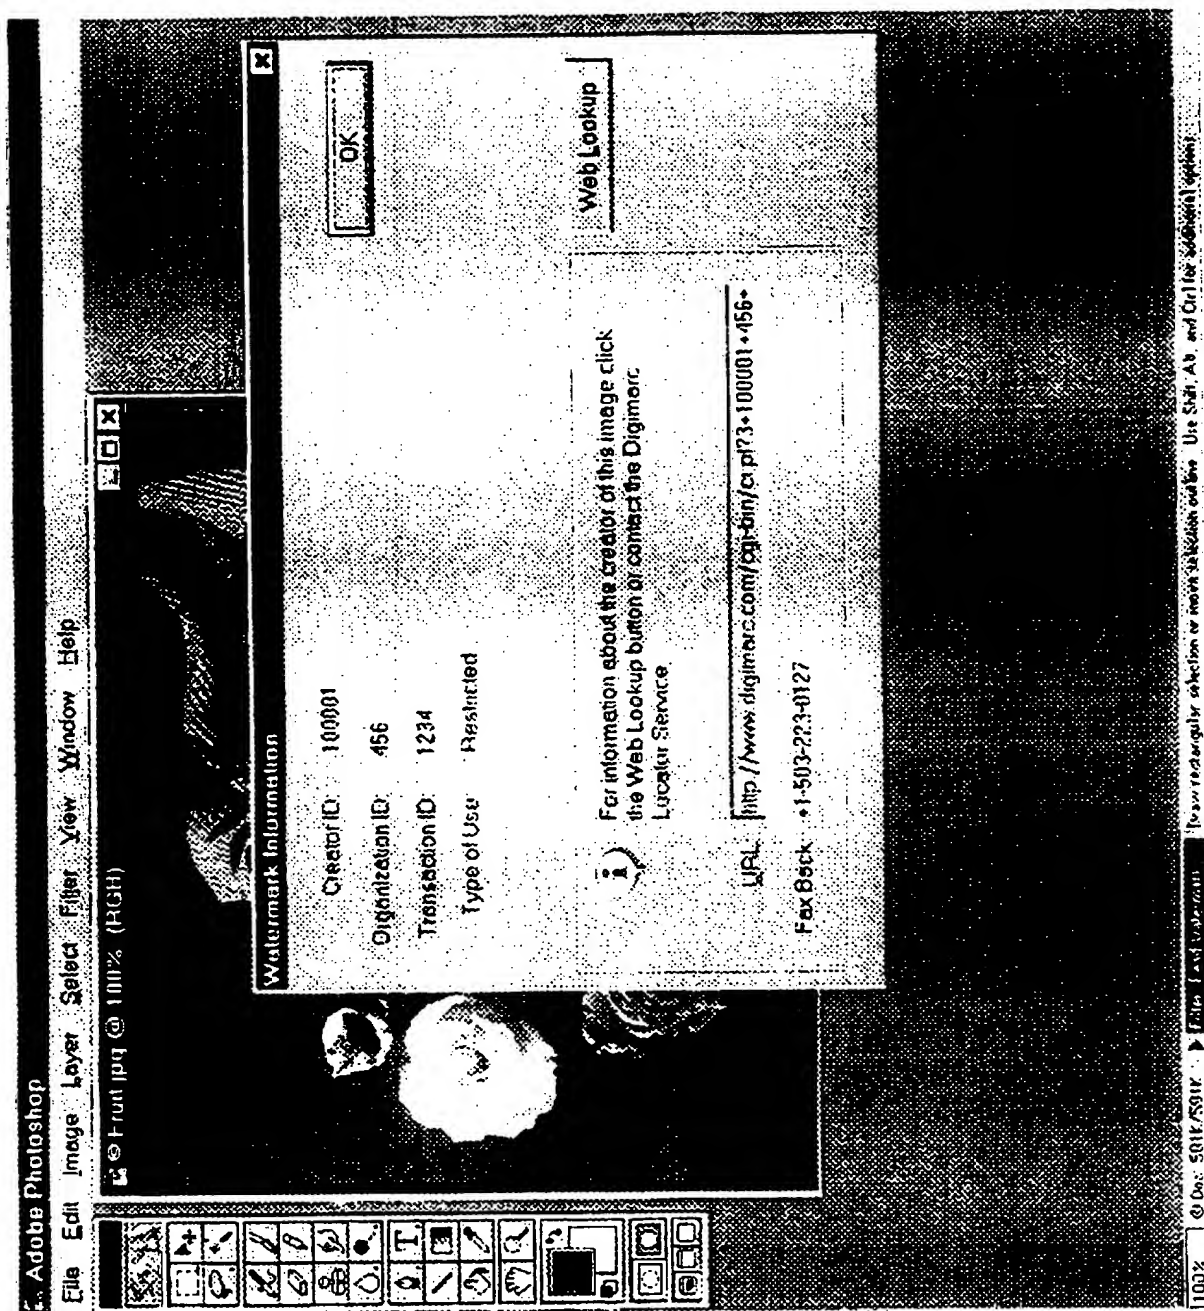


FIG. 56



FIG. 57






FIG. 59

51/55




FIG. 59




FIG. 60

POINTER, 1690 MAY ALSO BE STEGANOGRAPHICALLY EMBEDDED IN DATA STREAM IF DATA STREAM IS SO CAPABLE

RUDIMENTARY MASTER IDENTITY OBJECT AND SYSTEM


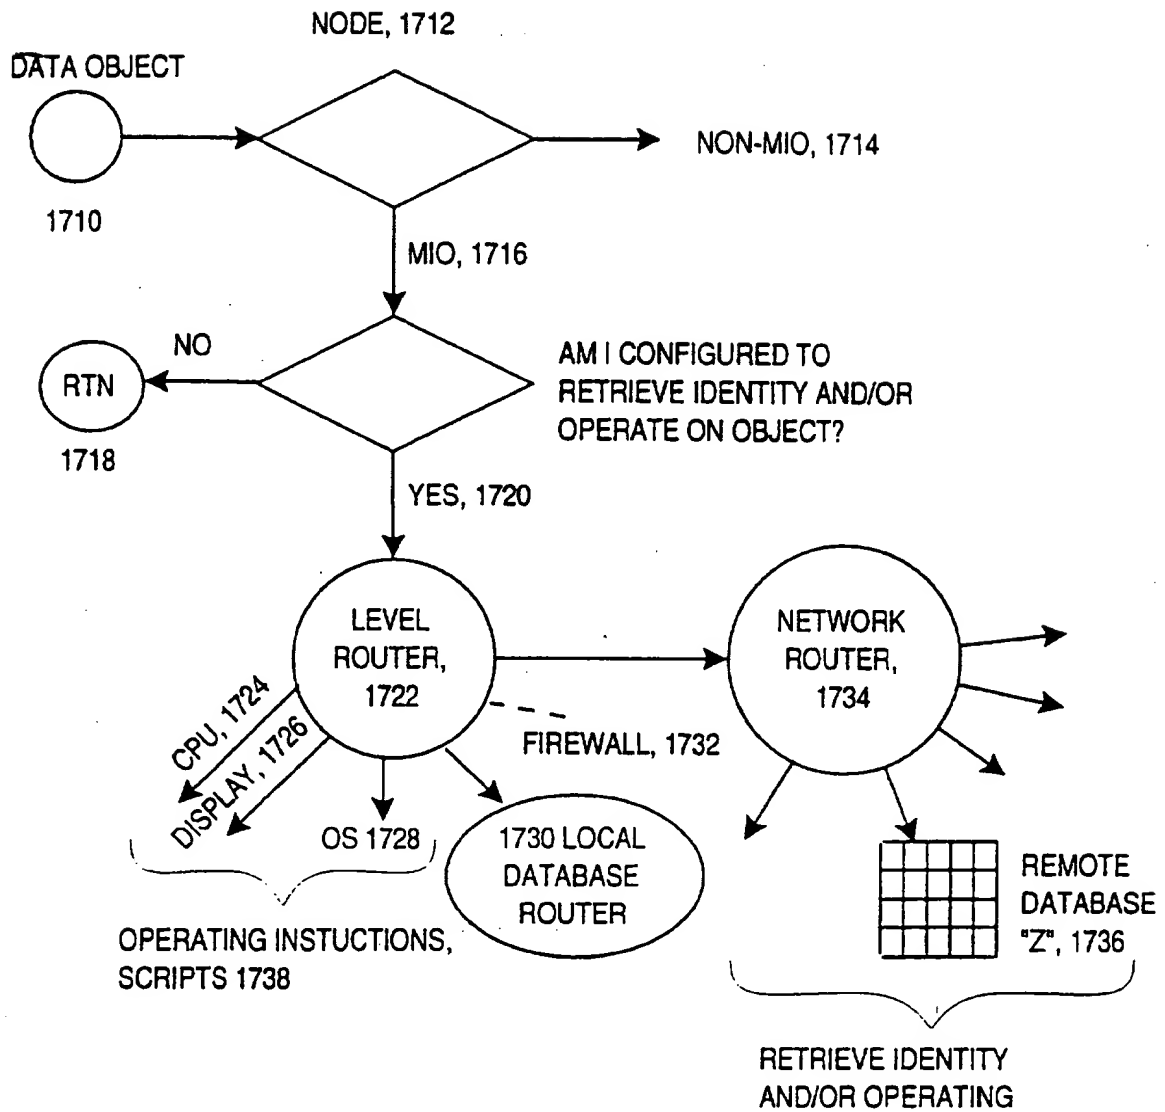


FIG. 61

53/55

**FIG. 62**

54/55

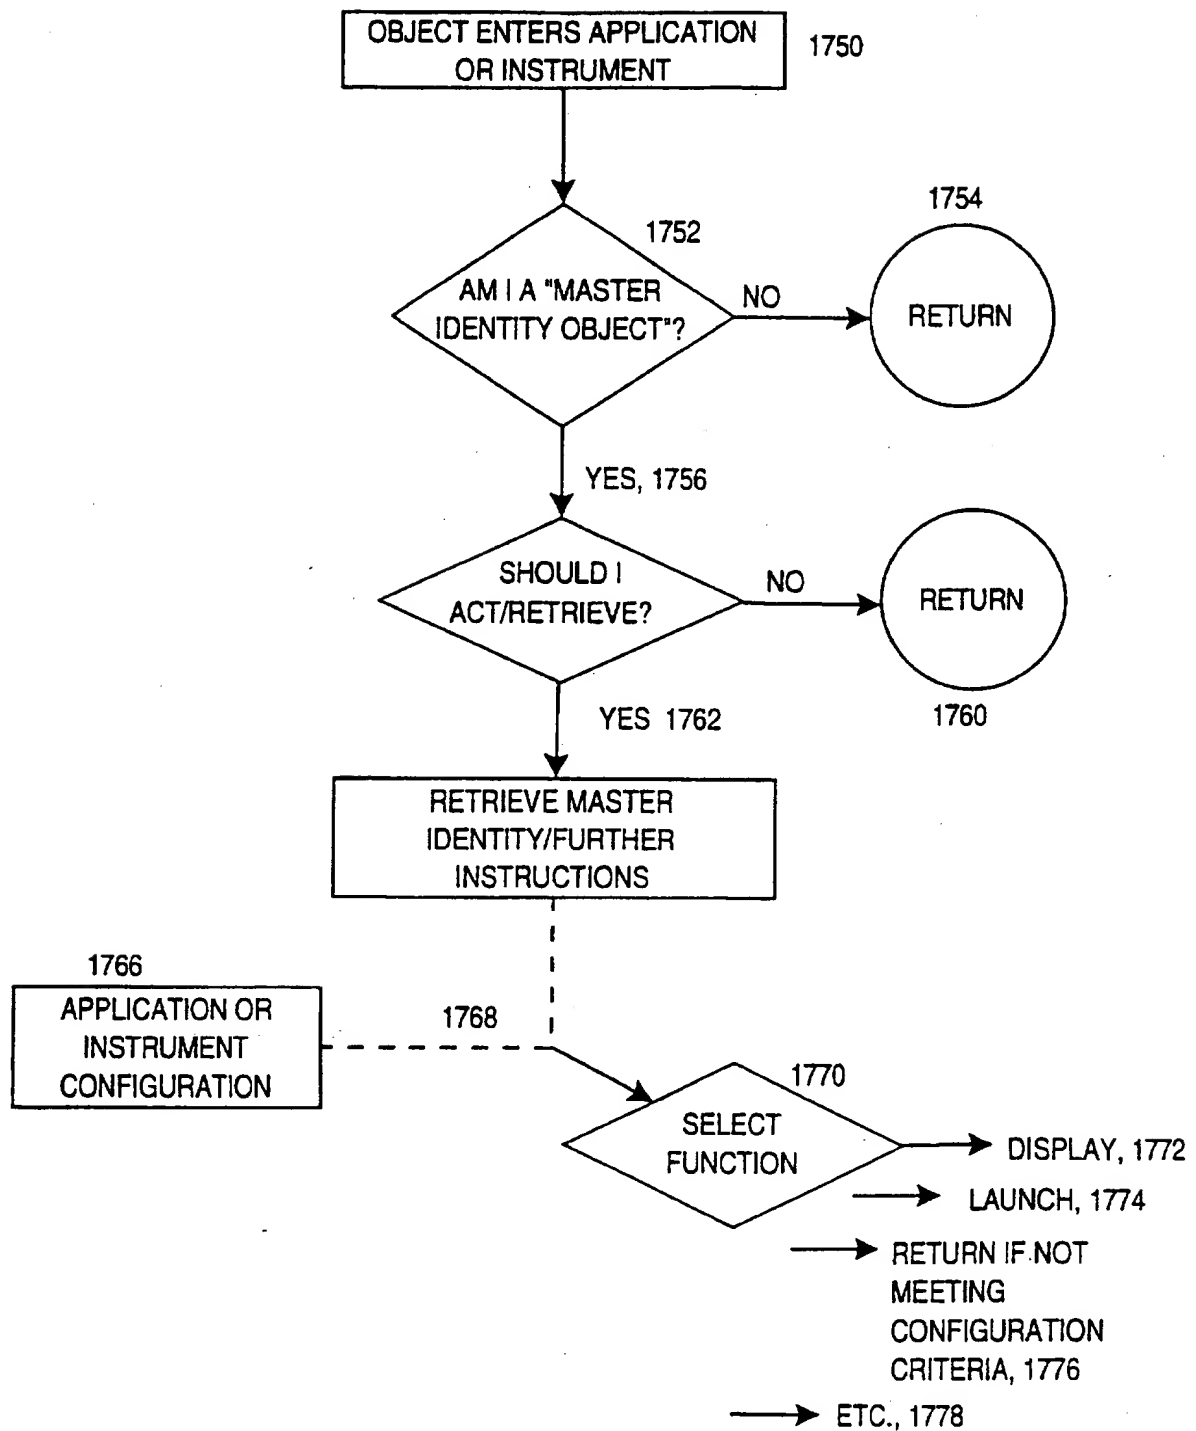
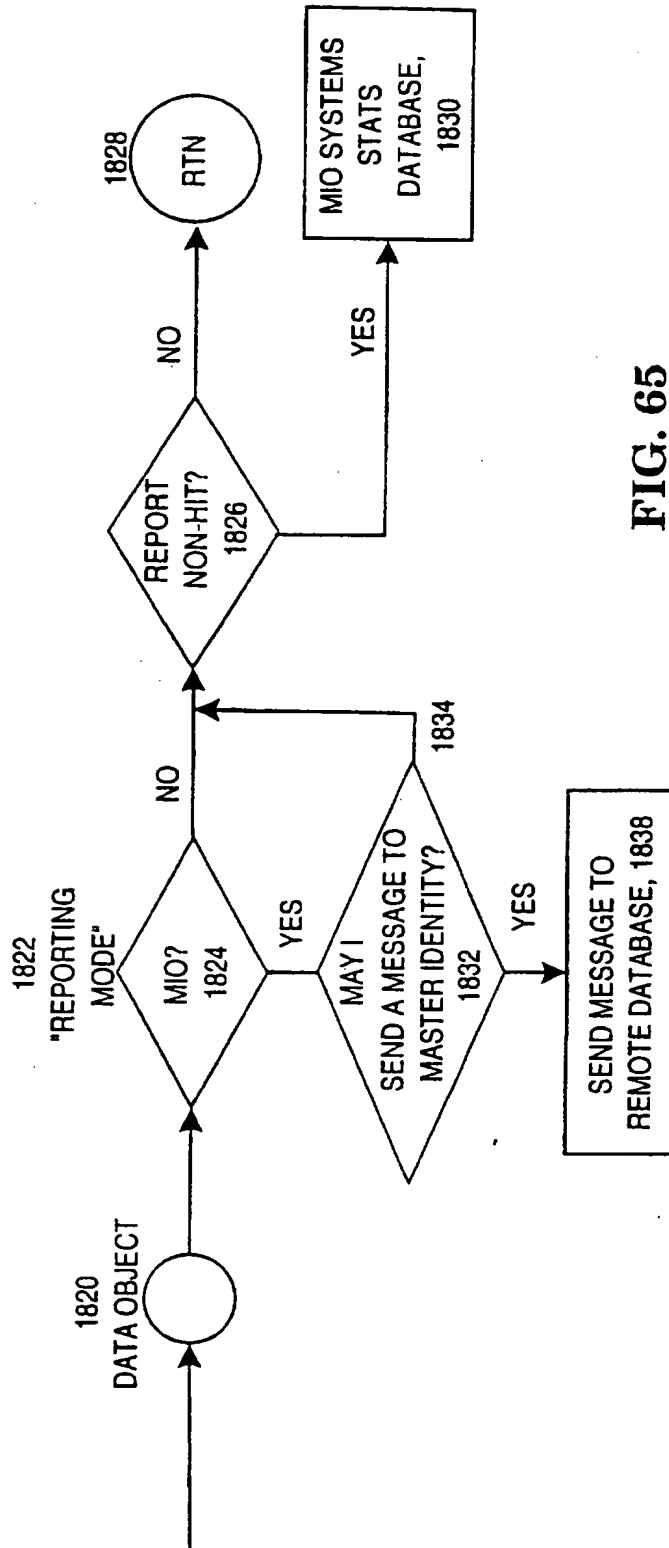



FIG. 63



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/08351

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06K 9/36

US CL : 382/232

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 382/232, 284; 395/135; 358/454,448; 380/4,23,25,54; 283/6,93,94; 358/142,124,450; 348/475; 235/494

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | EP 581 317 A2 (POWELL ET AL.) 02 FEBRUARY 1994, PAGES 1-6, FIGURES 1-5 | 1-5 |



Further documents are listed in the continuation of Box C.



See patent family annex.

| | | |
|--|------|--|
| * Special categories of cited documents: | * I | later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention |
| * A* document defining the general state of the art which is not considered to be of particular relevance | * X | document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone |
| * B* earlier document published on or after the international filing date | * Y | document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| * L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | * A* | document member of the same patent family |
| * O* document referring to an oral disclosure, use, exhibition or other means | | |
| * P* document published prior to the international filing date but later than the priority date claimed | | |

Date of the actual completion of the international search

17 SEPTEMBER 1997

Date of mailing of the international search report

31 OCT 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JOSE L. COUSO *Joni Hill*

Telephone No. (703) 305-4700

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US97/08351

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Please See Extra Sheet.

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-5

Remark on Protest

☐
☐

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.

BOX II. OBSERVATIONS WHERE UNITY OF INVENTION WAS LACKING

This ISA found multiple inventions as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single inventive concept under PCT Rule 13.1. In order for all inventions to be searched, the appropriate additional search fees must be paid.

Group I, claim(s) 1-5, drawn to a multi-computer system.

Group II, claim(s) 6-13, drawn to a sub-system for embedding information.

Group III, claim(s) 14, drawn to an image processing method.

The inventions listed as Groups I-III do not relate to a single inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: The claims of these three groups are directed to different inventions which are not so linked as to form a single general inventive concept since they recite different technical features. The inventions are not linked in operation and perform completely different operations, such as Group I, a multi-computer system, Group II, a system for embedding information, and Group III, an image processor.

Since applicant has not exercised the option, the International Preliminary Authority will establish the international preliminary examination report on those parts of the international application indicated below which are claimed first, namely Claims 1-5.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.